


NAS 1.21:519

DO NOT CIRCULATE

WITHDRAWN

Visualization Techniques in Space and Atmospheric Sciences

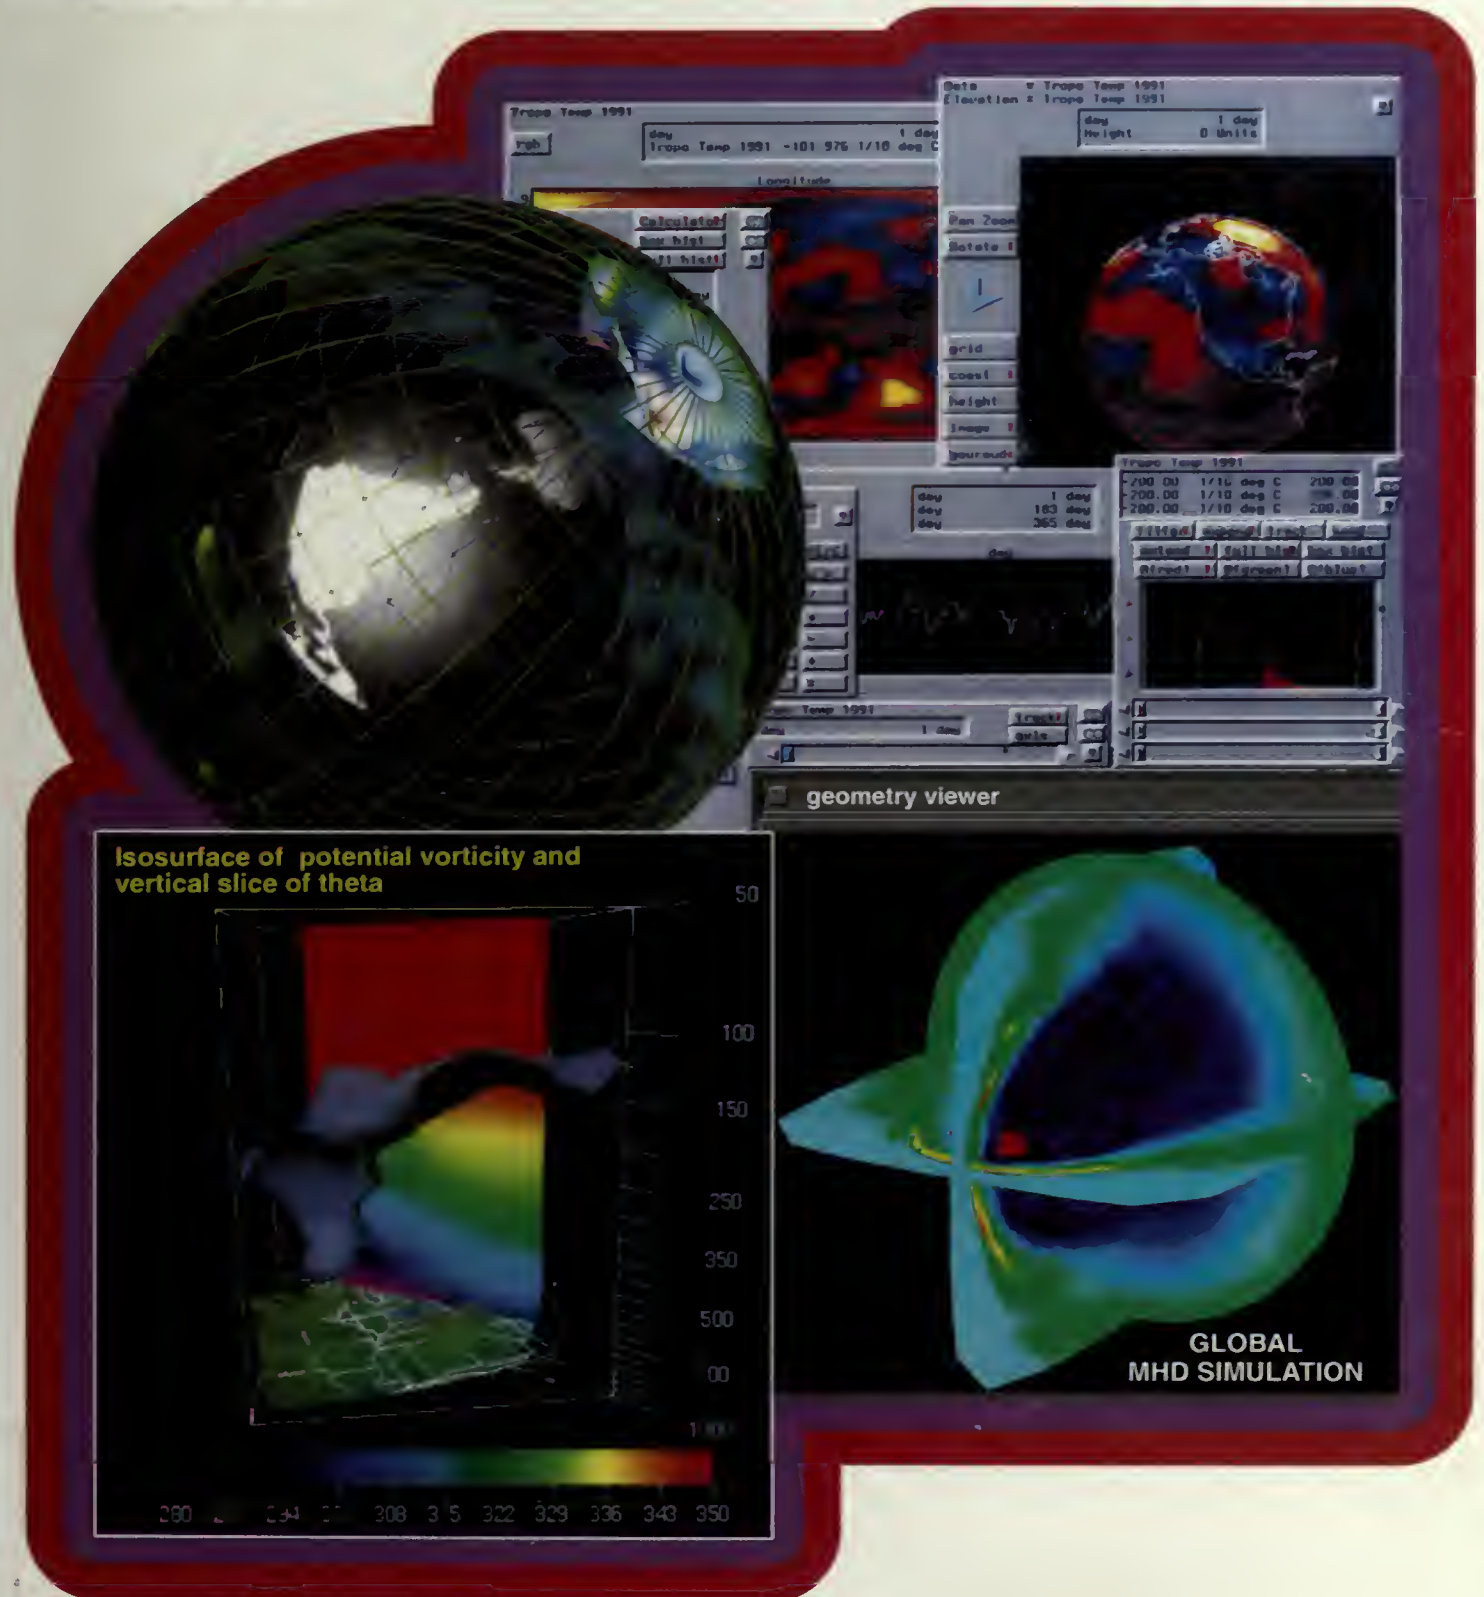
APACHE JUNCTION PUBLIC LIBRARY




3 9971 00211 4772

NOV 13 1995

DEPOSITIONAL ANALYSIS



Visualization Techniques in Space and Atmospheric Sciences



E.P. Szuszcwicz and J.H. Bredekamp
Editors



Digitized by the Internet Archive
in 2007 with funding from
Microsoft Corporation

Preface

The origin of this book can be traced to NASA's Applied Information Systems Research Program which was initiated to enhance space science productivity through the effective application of advanced computer science and technology. The program was developed because of the critical role that information systems and related technologies will play in meeting the scientific challenges of the 1990's and beyond. The program also has relevance as we expand our research agenda toward understanding the Earth as a system and eventually toward comprehending the origin of the universe.

Unprecedented volumes of data will be generated in pursuing these objectives, which will in turn require analysis and interpretation that will lead to meaningful scientific insight. Providing a widely distributed research community with the ability to access, manipulate, analyze, and visualize these complex, multidimensional data sets depends on a wide range of computer science and technology topics, and this book addresses many of the latest developments that service this critical need. Data storage and compression, data base management, computational

methods and algorithms, artificial intelligence, telecommunications, and high-resolution display are just a few of these topics.

The individual contributions are based on papers presented at a special session of the spring 1993 meeting of the American Geophysical Union (AGU). The session, "Advanced Data Handling and Visualization Tools for Space and Atmospheric Sciences," was a first for AGU in terms of its content and its use of on-line demonstrations of tools, technologies, and scientific insights. The need for interactivity, speed, user-friendliness, and extensibility is a unifying theme that the reader will find addressed throughout the papers.

The past 3 years have seen a major breakthrough, with newly developed systems beginning to appear on the desktops of working scientists and data analysts. We are on the threshold of a new and exciting era in science productivity as these capabilities evolve and attain broader application and interoperability.

Table of Contents

Preface	i
Table of Contents	iii
Acknowledgments	vii
I. OVERVIEWS	
Visualization: A Pathway to Enhanced Scientific Productivity in the Expanding Missions of Space and Earth Sciences	3
<i>E.P. Szuszczewicz</i>	
Data Handling and Visualization for NASA's Science Programs	13
<i>J.H. Bredekamp</i>	
II. SPACE AND UPPER ATMOSPHERIC SCIENCE APPLICATIONS	
Magnetic Field Experiment Data Analysis System	23
<i>D.B. Holland, L.J. Zanetti, L.L. Suther, T.A. Potemra, and B.J. Anderson</i>	
Scientific Visualization Tools for the ISTP Project: Mission Planning, Data Analysis and Model Interpretation	31
<i>M. Peredo, C.C. Goodrich, D. McNabb, R. Kulkarni, and J. Lyon</i>	
Visualization of the Electrostatic Potential Distribution in Both Polar Ionospheres Using Multiple Satellite Measurements	41
<i>M.R. Hairston, R.A. Heelis, and F.J. Rich</i>	
Application of Advanced Computing Techniques to the Analysis and Display of Space Science Measurements	51
<i>D.M. Klumppar, M.V. Lapolla, and B. Horblit</i>	
Data Management, Archiving, Visualization and Analysis of Space Physics Data	59
<i>C.T. Russell</i>	
Interplanetary Space Science Data Base and Access/Display Tool on the NSSDC Heliospheric CD-ROM	65
<i>N.E. Papitashvili, and J.H. King</i>	
SAVS: A Space and Atmospheric Visualization Science System	75
<i>E.P. Szuszczewicz, A. Mankofsky, P. Blanchard, C. Goodrich, D. McNabb, and D. Kamins</i>	

Visualization of International Solar-Terrestrial Physics Program (ISTP) Data	85
<i>R.L. Kessel, R.M. Candey, S.-Y. W. Hsieh, and S. Kayser</i>	
Visualization Software for a Global Three-Dimensional Upper Atmosphere Model	95
<i>B.T. Foster, R.G. Roble, and E.C. Ridley</i>	
A Prototype Upper Atmospheric Research Collaboratory (UARC)	105
<i>C.R. Clauer, D.E. Atkins, T.E. Weymouth, G.M. Olson, R. Niecejewski, T.A. Finholt, A. Prakash, C.E. Rasmussen, T. Killeen, T.J. Rosenberg, D. Detrick, J.D. Kelly, Y. Zambre, C. Heinselman, P. Stauning, E. Friis-Christensen, and S.B. Mende</i>	
Visualization Tools for the Processing of Airglow Data from RAIDS	113
<i>G.J. Miller, S. Thonnard, M. Picone, and C. Meesuk</i>	
Satellite Situation Center Data System for Magnetospheric Science Planning	121
<i>L. Aist-Sagara, J.F. Cooper, R.E. McGuire, R. Parthasarthy, and M. Peredo</i>	
Interactive Visualization of Numerical Simulation Results: A Tool for Mission Planning and Data Analysis ..	131
<i>J. Berchem, J. Raeder, R.J. Walker, and M. Ashour-Abdalla</i>	

III. LOWER ATMOSPHERIC AND EARTH SCIENCE APPLICATIONS

Using the Application Visualization System To View Haloe Three-Dimensional Satellite Data	145
<i>M. Luo, A.V.R. Schiano, J.M. Russell, III, L.L. Gordley, K.A. Stone, and R.J. Cicerone</i>	
A Versatile System for Processing Geostationary Satellite Data With Run-time Visualization Capability	155
<i>M. Landsfeld, C. Gautier, and T. Figel</i>	
Visualization of Stratospheric Ozone Depletion and the Polar Vortex	163
<i>L.A. Treinish</i>	
LinkWinds: A Visual Data Analysis System and its Application to the Atmospheric Ozone Depletion Problem	175
<i>A.S. Jacobson and A.L. Berkin</i>	
The Use of LinkWinds for the Validation and Analysis of 14 Years of Microwave Sounder Unit Daily Global Temperature Anomaly Data	185
<i>M.E. Botts and R.W. Spencer</i>	
Interactive Data Exploration and Particle Tracking for General Circulation Models	193
<i>R.I. Rosenbaum, R.L. Peskin, S.S. Walther, and H.P. Zinn</i>	
A Data Reduction, Management, and Analysis System for a 10-Terabyte Data Set	203
<i>R. DeMajistre and L. Suther</i>	
Geophysical Data Analysis and Visualization Using the Grid Analysis and Display System	209
<i>B.E. Doty and J.L. Kinter, III</i>	

IV. RELATED TOPICS

Making Sense From Space-Time Data in Laboratory Experiments on Space Plasma Processes	221
<i>W. Gekelman, J. Bamber, D. Leneman, S. Vincena, J. Maggs, and S. Rosenberg</i>	
Envision: An Interactive System for the Management and Visualization of Large Geophysical Data Sets	233
<i>K.R. Searight, D.P. Wojtowicz, J.E. Walsh, S. Pathi, K.P. Bowman, and R.B. Wilhelmson</i>	
Geophysical Phenomena Classification by Artificial Neural Networks	243
<i>M.P. Gough and J.R. Bruckner</i>	

DITDOS: A Set of Design Specifications for Distributed Data Inventories	253
<i>T.A. King, R.J. Walker, and S.P. Joy</i>	
Educational Software for the Visualization of Space Plasma Processes	263
<i>C.T. Russell, G. Le, J.G. Luhmann, and B. Littlefield</i>	
Data Base Management System and Display Software for the National Geophysical Data Center Geomagnetic CD-ROMs	271
<i>N.E. Papitashvili, V.O. Papitashvili, J.H. Allen, and L.D. Morris</i>	
Visual Interface for Space and Terrestrial Analysis	279
<i>E.G. Dombrowski, J.R. Williams, A.A. George, H.M. Heckathorn, and W.A. Snyder</i>	
PATHFINDER: Probing ATmospHeric Flows in an INtegrated and Distributed EnviRonment	289
<i>R.B. Wilhelmson, D.P. Wojtowicz, C. Shaw, J. Hagedorn, and S. Koch</i>	

Acknowledgments

We are pleased to acknowledge the untiring efforts of Ms. Charlene Urbanski whose role as Editorial Assistant covered everything from regular communications with the authors and reviewers to design of the cover. Her organization, efficiency, and amiable style were important elements in the overall effort. We also extend a special thank you to Ellwood Annaheim for his persistence and thorough attention to detail in preparing the manuscript for publication. We are also grateful for the people in NASA's Printing and Design Branch and their highly professional bookmaking. The editors and publication staff gratefully acknowledge the assistance of all reviewers who helped guide the coherence, clarity, and information content of the papers contained in this volume. Their task was not an easy one, because each review required an interdisciplinary perspective that covered computer system engineering, data handling, visualization technologies, and broad elements of space and atmospheric sciences. To those reviewers listed below, and to those who preferred to remain anonymous, we extend our heartfelt thanks.

*Acuna, M., Goddard Space Flight Center, Greenbelt, MD
Aiken, A. Goddard Space Flight Center, Greenbelt, MD
Allen, J. National Oceanic and Atmospheric Administration, Boulder, CO
Arnoldy, R., University of New Hampshire, Durham, NH
Bacon, D., Science Applications International Corporation, McLean, VA
Baker, D., Goddard Space Flight Center, Greenbelt, MD
Bayard-Cushing, J., Evergreen State College, Olympia, WA
Berchem, J., University of California, Los Angeles, CA
Berkin, A., Jet Propulsion Laboratory, Pasadena, CA
Blanchard, P., Science Applications International Corporation, McLean, VA
Botts, M., University of Alabama in Huntsville, Huntsville, AL
Boyd, D., University of Maryland, College Park, MD
Candey, B., Goddard Space Flight Center, Greenbelt, MD
Chappell, C.R., Marshall Space Flight Center, Huntsville, AL
Christensen, A., The Aerospace Corporation, Los Angeles, CA
Crain, D., Utah State University, Logan, UT
Crowley, G., John Hopkins University, Laurel, MD
Dombrowski, E., SFA, Inc., Landover, MD
Dryer, M., NOAA/Space Environmental Lab, Boulder, CO
Earle, G., Science Applications International Corporation, McLean, VA
Emery, B., High Altitude Observatory, Boulder, CO
Erickson, G., Boston University, Boston, MA*

- Esdale, P., Advanced Visual Systems, Inc., Waltham, MA
 Evans, D., NOAA/Space Environmental Lab, Boulder, CO
 Fejer, B., Utah State University, Logan, UT
 Flurchick, K., North Carolina Supercomputer Center, Research Triangle Park, NC
 Goodrich, C., University of Maryland, College Park, MD
 Grebowsky, J., Goddard Space Flight Center, Greenbelt, MD
 Heckman, G., NOAA/Space Environmental Laboratory, Boulder, CO
 Hedin, A., Goddard Space Flight Center, Greenbelt, MD
 Heelis, R., University of Texas at Dallas, Richardson, TX
 Hoffman, R., Goddard Space Flight Center, Greenbelt, MD
 Jacobson, A.S., Jet Propulsion Laboratory, Pasadena, CA
 Joselyn, J., National Oceanic and Atmospheric Administration, Boulder, CO
 Kelly, J., SRI International, Menlo Park, CA
 Kessel, R., Goddard Space Flight Center, Greenbelt, MD
 King, J., Goddard Space Flight Center, Greenbelt, MD
 Kinter, J., Center for Ocean-Land-Atmospheric Interactions, Calverton, MD
 Kress, M., CSI/CUNY, Staten Island, NY
 Kulkarni, R., University of Maryland, College Park, MD
 Lazarus, A., Massachusetts Institute of Technology, Cambridge, MA
 Lepping, R., Goddard Space Flight Center, Greenbelt, MD
 Luo, M., University of California, Irvine, CA
 McNabb, D., University of Maryland, College Park, MD
 Mendillo, M., Boston University, Boston, MA
 Moroh, M., Ken Wanderman and Associates, Shelter Island, NY
 Mucklow, G., NASA Headquarters, Washington, D.C.
 Neugebauer, M., Jet Propulsion Laboratory, Pasadena, CA
 Olson, J., University of Alaska, Fairbanks, AK
 Papitashvili, N., Goddard Space Flight Center, Greenbelt, MD
 Peredo, M., Hughes STX Corporation at NASA Goddard Space Flight Center, Greenbelt, MD
 Richmond, A., National Center for Atmospheric Research, Boulder, CO
 Rind, D., NASA/GISS, New York, NY
 Romick, G., John Hopkins University, Laurel, MD
 Salawitch, R., Harvard University, Cambridge, MA
 Sarma, R.A., Science Applications International Corporation, McLean, VA
 Schiano, A.V.R., University of California, Irvine, CA
 Searight, K., University of Illinois at Urbana-Champaign, Urbana, Ill.
 Snyder, W., Naval Research Laboratory, Washington, DC
 Solomon, S.C., University of Colorado, Boulder, CO
 Spiro, R., Rice University, Houston, TX
 Summers, M.E., Naval Research Laboratory, Washington, DC
 Suther, L.L., John Hopkins University/Applied Physics Lab, Laurel, MD
 Thorpe, S., North Carolina Supercomputer Center, Research Triangle Park, NC
 Tinsley, B., University of Texas at Dallas, Richardson, TX
 Torr, D., University of Alabama, Huntsville, AL
 Walker, D., Naval Research Laboratory, Washington, DC
 Walker, R., University of California, Los Angeles, CA
 Wicker, L., Texas A&M, College Station, TX
 Wilhelmson, R., NCSA, Urbana, IL
 Williamson, P.R., Stanford University, Stanford, CA
 Winningham, J.D., Southwest Research Institute, San Antonio, TX
 Wolf, R., Rice University, Houston, TX

Chapter I

OVERVIEWS

Visualization: A Pathway to Enhanced Scientific Productivity in the Expanding Missions of Space and Earth Sciences

E.P. SZUSZCZEWICZ

Laboratory for Atmospheric and Space Sciences
Science Applications International Corporation
McLean, VA

The movement toward the solution of problems involving large-scale system science, the ever-increasing capabilities of three-dimensional, time-dependent numerical models, and the enhanced capabilities of "in situ" and remote sensing instruments bring a new era of scientific endeavor that requires an important change in our approach to mission planning and the task of data reduction and analysis. Visualization is at the heart of the requirements for a much-needed enhancement in scientific productivity as we face these new challenges. This article draws a perspective on the problem as it crosses discipline boundaries from solar physics to atmospheric and ocean sciences. It also attempts to introduce visualization as a new approach to scientific discovery and a tool which expedites and improves our insight into physically complex problems. A set of simple illustrations demonstrates a number of visualization techniques and the discussion emphasizes the trial-and-error and search-and-discover modes that are necessary for the techniques to reach their full potential. Further discussions also point to the importance of integrating data access, management, mathematical operations, and visualization into a single system. Some of the more recent developments in this area are reviewed.

1. THE GRAND CHALLENGE OF LARGE-SCALE SYSTEM SCIENCE

Space and Earth Science is shifting in emphasis from single exploratory investigations to concentration on collaborative missions of large-scale system phenomena producing unprecedented volumes of data with unmatched temporal, spectral, and spatial resolution. The multi-parameter, multi-platform data bases will challenge our intellectual capabilities to absorb, synthesize, and effectively analyze the measurements, as well as impose new demands on software and hardware systems as we know them today.

The large-scale numerical codes intended to model, interactively analyze, and ultimately predict the observed phenomena are of an equal challenge, requiring their implementation on massively parallel computers with performance characteristics approaching the teraflop range. These codes, whether they address individual disciplines like solar, heliospheric, or atmospheric physics, or the coupling of solar-terrestrial processes that control our weather, have commonality in their complexities. They all involve three-dimensional time-dependent geometries, large variations in temporal and spatial scales, complex physical interactions, and most often non-linear dynamics.

If we are to meet the challenge of the National Academy of Sciences to be able to predict the Earth's environment in the context of its position in the solar system and in the context of global change, we must: (1) develop the systems which can efficiently gather and process the data, (2) access, compare and "tune" the appropriate models, and (3) integrate the results into a predictive capability. Here the integration of human intelligence with hardware and software systems, in addition to

appropriately designed visualization tools, plays a pivotal role if these ambitious goals are to be achieved and a predictive capability developed.

In the domain of solar and heliospheric physics, we expect three-dimensional magnetohydrodynamic (MHD) models to be among the basic tools for a coherent hierarchical approach to the study of solar processes and their control of the global heliosphere from its coronal origins to its interactions with the Earth's bow shock and the interstellar medium. The origins of the solar wind as well as coronal mass ejections (CMEs) must be modeled to help interpret ground-based and spacecraft data, such as that from the Solar Heliospheric Observatory (SOHO). Similarly, the interaction of solar wind streams, the propagation of CMEs, and the structure of the heliosphere out of the ecliptic must be modeled and compared with data from the Ulysses spacecraft. And the distant interaction of the solar wind with the local interstellar medium, leading to the formation of the solar wind termination shock, the heliopause, and the heliospheric bow shock, must be modeled to help interpret data from the Pioneer and Voyager spacecraft now reaching towards the edge of the solar system.

The work of the solar and heliospheric communities includes the requirements of magnetospheric physicists who will be more intent in looking out into the interplanetary medium and dynamic solar phenomenologies to understand the cause-effect relationships in their discipline area of geospace. Upper atmospheric and thermospheric physicists will continue to be more keenly aware of the need for an understanding of ionospheric coupling mechanisms and the controls from the magnetosphere and the heliosphere above, and the stratosphere below. In addition, scientists studying the Earth's ionospheric-thermospheric-mesospheric system will be looking to Mars, Venus, and beyond for comparative studies of large-scale planetary environments that not only form an intrinsic element in solar-system exploration, but provide unique opportunities to determine the limits of Earth-based concepts and extend our fundamental understanding of the processes under study.

Atmospheric and ocean scientists will be looking for improved global-scale measurements of sea state and surface winds to better understand air-sea interactions, heat transfer mechanisms, weather,

and climate. And the hydrologic cycle will be an experimental and theoretical focus because of its important role in creating and modulating the various climatic zones on the Earth. The global circulation models will therefore require the integration of large volumes of environmental data, including rain-rate and rain-distribution data from ground-based radars, rain gauges, and spaceborne sensors like those included in the Tropical Rainfall Measurement Mission (TRMM).

The increased focus on large-scale system phenomena, the cross-disciplinary nature of many investigations, and projections of increased volumes of "in situ" and remote sensing data are all elements in the coming decade of Earth and Space Sciences. These projects range from new astronomical observatories such as the Space Telescope, planetary orbiters such as Galileo and Cassini, the multi-spacecraft International Solar Terrestrial Physics (ISTP) and TIMED missions, and the EOS mission to planet Earth. ISTP alone is expected to accumulate almost a gigabit of data per second.

To be effective in their work scientists will require a data analysis environment with sophisticated interactive graphics tools based on advanced visualization techniques. Such techniques represent one of the most effective ways to "compress" complex data into a visually organized form which is optimized for analysis and interpretation. The scientists also require easy-to-use mathematical and image processing tools. In addition, they will need tools to assist in obtaining data from remote archives as well as to access the results of empirical and numerical models to correlate with the data and assist in analysis and interpretation. To be most effective, these tools must be integrated into a single user-friendly system.

2. VISUALIZATION: THE PATHWAY TO ENHANCED SCIENTIFIC PRODUCTIVITY

Visualization is at the heart of our requirements for the much needed enhancement in scientific productivity as we face the grand challenge of large-scale system phenomena in space and Earth sciences. This is not to underplay the important roles of remote data access or easy-to-use mathematical tools, because they are the fundamental elements

with which the visualization process begins.

In itself, scientific visualization is a rapidly evolving field. But the reality is that we have been applying visualization techniques in one form or another during our entire scientific careers. What has changed has been the magnitude of the problems we are addressing, the sensitivity and dimensionality of our instruments, the volume of our data base, and the hardware, software, and computational tools to interpret and assimilate the results. Visualization is not new. We are simply in a new era that requires a dramatic change in our approach and our attitudes toward data handling and analysis. This attitude must recognize that visualization is not a “pretty picture,” but rather a process of discovery that links the knowledge of the discipline with the knowledge of visualization tools. In the past, what may have been called visualization was really “data display,” that is, the graphical rendering of the data when we knew what it meant or when we knew what we wanted it to say. On the other hand, the “visualization process,” as we have come to know it today, is the intersection of intellectual pursuit with a rendering of the data or model results which leads to discovery. The rendering involves the use of color, intensity, transparency, texture, animation and many other techniques which can convey a tremendous amount of information in a single image and in a short period of time. The overlay of images with varying degrees of transparency and texture, or the superposition of scalar and vector fields, or a sequence of images (i.e., animation) can yield more information than the sum of the parts. This allows our intellectual capabilities to parallel process the results as no machine can do, and couples the knowledge of our discipline with the knowledge of the visualization tools.

Emphasizing its functional aspects, visualization can be thought of as having two major components: exploration and presentation [Keller and Keller, 1993]. In the exploration mode, we search the data or model results for new relationships or insights into complex phenomena. Often this means trial-and-error representations of the data and/or model results, requiring interactive adjustments of the data or the image and the direct participation of the scientist’s knowledge of the discipline. It is in this mode that we emphasize visualization as a

“process” rather than the graphical end-product of the “presentation” mode in which the arrangement and display of the data is for publication and the benefit of others.

It must be emphasized that the exploitation of visualization technologies and the enhancement in scientific productivity require a somewhat non-traditional approach to handling and rendering the data and model results. Unfortunately, many scientists have been conditioned to regard data and model outputs as entities with inviolate properties that ultimately dictate the use of standard displays and presentation formats for the data. Such rigid thinking precludes the exploitation of visualization techniques and undermines the exploration phase where trial and error in the coupling of knowledge of the discipline with the knowledge of the visualization tools is the pathway to discovery.

3. ILLUSTRATIONS

Recognizing that visualization is a process involving tools that include the use of color, transparency, texture, etc., along with rotation, pan, zoom, and animation [see e.g. Friedhoff and Benzon, 1989; and Nielson et al., 1990], this section presents simple illustrations of several visualization techniques to highlight their utility in unfolding complex data, in developing quick insights into a problem, and in conveying large amounts of information in short periods of time.

Color vs. isolines and multi-parameter overlays. Color is one of the simplest tools for visualizing data and it finds extended use in identifying morphologies in a scalar field. Isolines can do this as well, but color can identify morphologies when isolines fail. This is illustrated most simply with the blue-to-red color rendering of the scalar field of numbers ranging from 0 (blue) to 9 (red) in Figure 1. This might be a field of temperature measurements over a land mass or over the ocean. It might also represent plasma density distributions under disturbed conditions in the polar cap ionosphere. In any case, the visualization task must unfold unique morphological features (i.e., concentrations, rarefactions, etc.). Isolines would fail in this application, since they would yield a tangled net of “spaghetti” that provides no simple organization of the data. On the other hand, the use of color immedi-

ately identifies a “hot spot” of 9s near the center of the image in Figure 1.

Color has found great utility in a number of disciplines where there is the need for multi-




Figure 1. A color rendering of a sample two-dimensional array of numbers (0 [blue] through 9 [red]) mimicking any scalar field of measurements where morphologies or topologies need to be determined. The color representation of the numerical field clearly reveals a concentration of high-scale measurements (9s) near the center of the image, and a random distribution elsewhere [from A. Mankofsky, private communication (1993)].

parameter overlays in endeavors focused on understanding a phenomenon rather than a single data set. A typical application is in thermospheric physics, where two dimensional displays in latitude and longitude can, for example, depict atomic oxygen densities in color, molecular nitrogen densities as isolines, and winds with vectors (see Figure 2). Such an image is effective, but only at a given altitude. Understanding the full three-dimensional, time-dependent aspect of thermospheric circulation and its coupling to the global-scale ionosphere under magnetic storm conditions requires more advanced visualization tools than those demonstrated in Figure 2. Such tools include isosurfaces, streamlines, transparency, texturing, and animation, all of which are available in a number of commercial software packages (e.g., AVS from AVS, Inc., Data Explorer from IBM, etc.), and some of which have been implemented already in developmental applications [e.g., Foster et al., 1994; Twiddy et al., 1994; and Gekelman, 1994].

Isosurfaces and the need for a non-traditional approach to data presentation. The exploratory




Figure 2. Illustration of color overlays to superimpose geophysical parameter sets in a study of the Earth's thermosphere. Image shows the global distributions of winds (white vectors), molecular nitrogen densities (color-coded isolines), and atomic oxygen densities (full color-coded cut-plane) at 300 km as specified by the MSIS model [Hedin, 1991].




Figure 3. Illustration of a non-traditional approach to rendering data and modeling results. Visualizing the Earth's geomagnetic field with an isosurface application to the IGRF model immediately reveals the phenomenological domain called the South Atlantic Anomaly. See text for discussion.

nature of the visualization process requires that scientists free themselves from rigid historical views on the display and presentation of their data. When this is done the opportunity presents itself for entirely new perspectives on the data and associated insights into the prevailing physical principles. This is illustrated in Figure 3 with a non-traditional approach to the representation of the Earth's geomagnetic field using the International Geomagnetic Reference Field model [IGRF, available from the National Geophysical Data Center, Boulder, CO].

Typically, there are two standard representations of the Earth's field. The first involves two-dimensional isocontours (i.e., isolines) of B-field values at a fixed altitude. The second is a fieldline representation that demonstrates the field's dipolar nature. A non-traditional approach is to view the geomagnetic field with a three-dimensional isosurface of sequential B-field values. One such image for the $B=0.5$ and 0.265 Gauss isosurfaces is presented in Figure 3, which shows the $B=0.5$ surface (yellow) confined to the polar regions. The $B=0.265$ surface is more global, with the image showing it furthest from the Earth near the north and south polar regions, closest to the Earth at mid- and equatorial latitudes, and intersecting the Earth in the South American/South Atlantic region. This intersection locates the well-known South Atlantic

Anomaly where the field at the surface of the Earth has anomalously small values and where, as a consequence, energetic particles are known to penetrate to very low altitudes [Ratcliffe, 1972]. This picture presents an intuitive insight into the causality, recognizing that the isosurface represents an approximation to the magnetic pressure ($B^2/8\pi$) and that energetic particles have higher probabilities of atmospheric entry in the South Atlantic "pressure well" that exists naturally in the Earth's B-field topology. One can only speculate on the impact that such a visualization tool may have had in early studies of the Earth's field and its control of charged particle dynamics.

The visualization process in data reduction and model-measurement comparisons. The products of visualization tools are not always representations of data and model runs. They often provide tests of model or algorithm integrity and time-saving views of an experiment scenario that facilitate determination of the shortest path in the process of data reduction and analysis. These aspects of visualization are illustrated in an application to a rocket-borne chemical release experiment in the NASA/CRRES program that involved "in situ" and ground-based systems and aspects of coupled phenomena along geomagnetic field lines.

The first panel of Figure 4 shows the rocket trajectory along with an array of B-field vectors determined by the IGRF. The obvious discontinuity in the B-field representation revealed a problem in an interpolation module that may have gone undetected without this simple visual aid.

The second panel of Figure 4 illustrates the application of visualization tools to help understand the actual configuration of the experiment. In the simplest description of the experiment scenario, a Ba/Li mixture in a sealed canister was separated from the rocket's diagnostics payload and the vaporized gases were released on the upleg portion of the trajectory at an altitude near 280 km. The colored disk in the panel is a color-coded model depiction of the cloud's ion density in the plane perpendicular to the cloud's bulk velocity at a very early phase of the expansion process. The red line is the suborbital trajectory, and the long and short vectors, respectively, are the B-field and payload velocity at the point of the release. The cone-shaped object in the figure is the projection of a ground-



Figure 4. Illustration of visualization as an integral process of understanding an experiment configuration and optimizing the approach to data reduction and analysis. See text for details.

based hf heater beam intended to intercept the chemical cloud at its ionospheric altitude, excite the electrons, and induce enhanced optical emissions and plasma instability processes.

The initial process of understanding the actual configuration of the experiment involved the following questions:

1. What were the magnitude and directions of the cloud's bulk velocity relative to the geomagnetic field?
2. How well was the chemical release coupled along the B-field to the diagnostic payload?
3. How effectively did the heater beam intercept the cloud?
4. Was it possible to observe expanding cloud ions on the downleg portion of the trajectory?

The first three questions were answered using rotation and zoom tools, with some of the image products presented in the middle and right-hand panels of Figure 4. The last question was answered with yet another rotation that projected an image perpendicular to the plane of the trajectory, with the downleg portion of the trajectory in the foreground. That image, presented in the right-hand panel of Figure 4, shows that the B-field at the release site did not intersect the downleg trajectory. A simple hand-scaling of altitude and range showed that the

fluxtube passing through the release missed the downleg trajectory by 13 km, a sufficiently great distance to minimize interests in searching the data for possible late-time ion signatures of the expansion process on the downleg trajectory.

The entire exercise of answering the questions listed above required less than five minutes. Without the visualization tools, the answers would have required the development of a number of mathematical algorithms and a sequence of stack plots. That exercise could have taken 3 to 5 days, depending on the proficiency of the programmer.

This illustration is one of many that could have been presented to highlight visualization as a time-saving and insightful element in the process of data reduction and analysis. This is especially true in situations that require co-registration of spaceborne "in situ" and remote sensing techniques supported by ground-based diagnostics. For example, the NASA TIMED mission will rely heavily on spaceborne and ground-based remote sensing techniques in its attempt to understand the coupling of the ionospheric, thermospheric and mesospheric domains and the influence of magnetospheric processes. The three-dimensional mapping of gravity wave influences, dynamic changes in neutral composition, field-line integrated conductivities, and penetrating electric field events are among the many

issues that will need to be addressed. This will require considerable care in space and time correlations along with co-registration of instrument fields-of-view with coupled regions of the mesosphere, the lower thermosphere, and the D, E, F₁, and F₂ regions of the ionosphere.

On the rendering of three-dimensional processes with isosurfaces, cut-planes, and transparencies.

The problem of co-registration and space-time correlations among measurement platforms and instrument types can be particularly acute in investigations of the solar wind-magnetosphere-ionosphere (SW-M-I) system. Such investigations involve an incredible volume of space with prevailing phenomena that span the spatial spectrum from meters to several Earth radii. The coupling mechanisms within the SW-M-I system are at the center of the NASA ISTP and the NSF GEM programs, with a focus on the flow of energy and momentum into and through the system. The difficulties include the complexities of the interacting phenomena, the large set of controlling parameters, and the coupled domains of varying dimensions.

The utility of several visualization tools in the study of the SW-M-I is illustrated here through applications to the MHD model of J. Lyon [Fedder and Lyon, 1987] in Figure 5. This figure, the product of the SAVS data acquisition and visualization system [Szuszczewicz, et al., 1994], has been created to address a Gedanken Experiment that attempts to determine the flow of momentum into specified regions of space under steady-state and dynamic environments. The upper panel presents the following elements: (1) the inner boundary of the computational domain at 3.5 R_E in red, (2) open and closed field lines in white, and (3) a pair of transparent isosurfaces (i.e., isobars of constant plasma pressure) with a color scale that represents the magnetic pressure term $B^2/8\pi$ everywhere on the isobaric surface. (Transparency is used to show the B-field structure within the isobaric surface).

The lower panel is a two-angle rotation (i.e., pitch and roll) of the top image with the transparency of the isobaric surfaces eliminated. The cut-plane, which is perpendicular to the plane of the B-field lines, presents: (1) bulk plasma velocity vectors in white, (2) the plasma density in a slightly transparent color code, and (3) the “spider web” computational grid of the MHD model which

provides varying degrees of resolution in the computational domain. This image sets the stage for an improved understanding of the topologies and morphologies in the overall system and for the calculation of momentum and energy flows into and out of critical domains. In this case the Gedanken Experiment is concentrated on the isobaric surfaces whose topology we see is controlled by the lowest-latitude open field lines in the nightside hemisphere.

These are just two of many possible images that can be at the researcher’s disposal. They clearly display the morphological relationships among the computational grid, the field configuration, the kinetic and magnetic pressures, the plasma density, and the bulk flow velocities—all elements critical to the calculation of energy and mass transfer. Without such images, insight and understanding of the problem would be impaired.

4. THE NEED FOR AN INTEGRATED SYSTEM

Recognizing that visualization is a process that enhances insight into the “reality” of physically complex problems is just the beginning for enhanced scientific productivity in the expanding missions of space and Earth sciences. Visualization is indeed a condition without which we cannot effectively address the grand challenges of the coming decade. Scientific productivity demands an interactive data analysis environment with advanced visualization tools, because such an approach represents one of the most effective ways to compress complex data into a visually-organized form that is optimized for analysis and interpretation.

While it is generally agreed that easy-to-use visualization and data handling tools are vital to efficient and insightful progress in science, the applications development environment has been characterized as fragmented, with little overall direction or coordination [Botts, 1992]. It has been argued that the primary bottleneck is the lack of adequate software which allows the scientist to take advantage of today’s computing power and interactively visualize and analyze his/her data within the complex computing environment. The end result is that the largest fraction of the scientific community is not using the visualization capabilities available today because (according to Botts): (1) the tools are

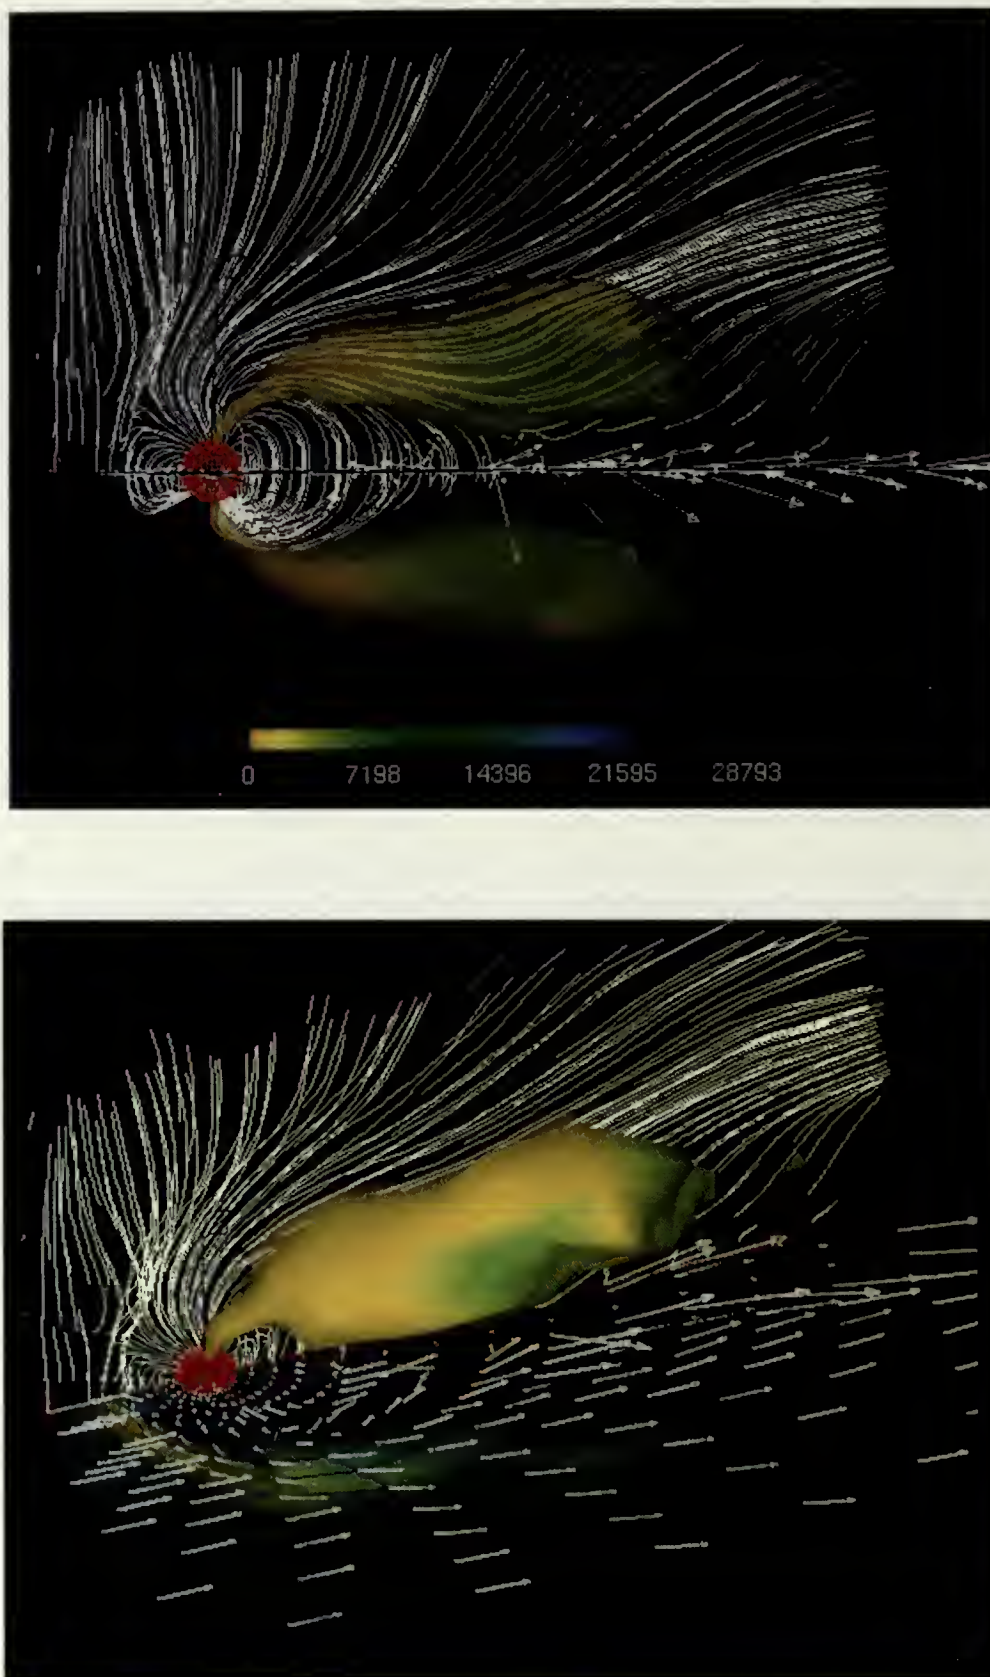


Figure 5. Illustration of imaging and visualization perspectives in addressing magnetospheric problems. Based on the MHD model of Fedder and Lyon (1987), this figure includes representations of the computational grid along with plasma and magnetic field pressure, bulk velocities, plasma densities, and open/closed field lines. See text for complete discussion.

not extensible or are too inflexible, (2) it is too difficult to input data, (3) the tools do not adequately link visualization and analysis, and (4) the tools do not meet scientific expectations. It is also agreed throughout the scientific community and effectively argued by Botts that the main general areas needing consideration include:

1. integration of visualization with data management and analysis;
2. a framework of flexibility without complexity; and
3. the development of an architecture focused on the scientist as the end user.

This requires that the tools meet the actual needs of the scientist, be simple and intuitive to use, and be logical to the scientist rather than to a computer specialist.

Many of these criticisms are currently being addressed within NASA's Advanced Information Systems Research Program in the Office of Space Science, with progress having been reported recently in a special session of the American Geophysical Union [EOS 74, No. 16, April 20, 1993/Supplemental] and its published proceedings [Szuszczewicz and Bredekamp, 1994; this volume]. As an illustration, the works of Berchem et al. (1994), Jacobsen and Berkin (1994), Szuszczewicz et al. (1994), Peredo et al. (1994), Russell (1994), and Searight et al. (1994) are developing integrated approaches to the access, display, and analysis of multivariate multidisciplinary data sets. In a number of cases (e.g., Berchem et al., Szuszczewicz et al., and Peredo et al.), the integrated approach includes multifunctional data readers, access to large-scale numerical models, and the superposition of satellite orbit tracks for model-measurement comparisons. The efforts of Jacobsen and Berkin (1994) and Szuszczewicz et al. (1994) emphasize ease, functionality, and extensibility with simple push-button interfaces to data acquisition, mathematical processing, and visualization tools in an interactive environment that addresses a broad spectrum of space and Earth sciences.

For these systems to reach their full potential there is the intrinsic need for scientists to open their minds to the broad spectrum of capabilities imbedded in the visualization tools that are being made available. As discussed earlier, many scientists have been conditioned to regard data and model outputs

as entities with inviolate properties that dictate the use of standard displays. This attitude must be abandoned for the full potential of visualization to be realized. It is a rapidly developing field, that holds great promise. But like any resource it must be effectively mined.

REFERENCES

- Berchem, J., J. Raeder, R.J. Walker, and M. Ashour-Abdalla, Interactive Visualization of Numerical Simulation Results: A Tool for Mission Planning and Data Analysis, in *Visualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszczewicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Botts, M.E., The State of Scientific Visualization with Regard to the NASA EOS Mission to Planet Earth, Report to NASA Headquarters, July 1992.
- Fedder, J.A. and J.G. Lyon, The solar wind-magnetosphere-ionosphere current-voltage relationship, *Geophys. Res. Lett.*, 14, 880, 1987.
- Foster, B.T., R.G. Roble, and F.A. Marcos, Time-Dependent Animation Software for Analysis of NCAR Thermosphere/Ionosphere General Circulation Model (TIGCM) Results, and Comparison with Satellite Data, in *Visualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszczewicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Friedhoff, R.M. and W. Benzon, *The Second Computer Revolution: Visualization*, Harry N. Abrams, Inc., New York, 1989.
- Gekelman, W., Making Sense from Space-Time Data in Laboratory Experiments on Space Plasma Processes, in *Visualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszczewicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Hedin, A.E., M.A. Biondi, R.G. Burnside, G. Hernandez, R.M. Johnson, T.L. Killeen, C. Mazaudier, J.W. Meriwether, J.E. Salah, R.J. Sica, R.W. Smith, N.W. Spencer, V.B. Wickwar, and T.S. Viridi, Revised Global Model of Thermosphere Winds Using Satellite and Ground-Based Observations, *J. Geophys. Res.*, 96, 7657, 1991.

- Jacobson, A.S. and A.L. Berkin, LinkWinds, A Visual Data Analysis System and its Application to the Atmospheric Ozone Depletion Problem, in *DVisualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszcwicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Keller, P.R., and M.M. Keller, *Visual Clues: Practical Data Visualization*, IEEE Computer Society Press, Manning Publ. Col., Greenwich CT, 1993.
- Nielson, G.M., B. Shriver, L.J. Rosenblum, editors, *Visualization in Scientific Computing*, IEEE Computer Society Press, Los Alamitos, CA, 1990.
- Peredo, M., C.C. Goodrich, D. McNabb, R. Kulkarni, and J. Lyon, Scientific Visualization Tools for the ISTP Project: Mission Planning, Data Analysis and Model Interpretation, in *Visualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszcwicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Ratcliffe, J.A., *An Introduction to the Ionosphere and the Magnetosphere*, Cambridge University Press, 1992.
- Russell, C.T., Data Management, Archiving, Visualization and Analysis of Space Physics Data, in *Visualization Techniques in Space and Atmospheric Sciences*, edited by E.P. Szuszcwicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Searight, K.R., S. Pathi, K.P. Bowman, D.P. Wojtowicz, R.B. Wilhelmson, and J.E. Walsh, Envision: An Interactive System for the Management and Visualization of Large Geophysical Data Sets, in *Developments and Applications of Data Handling and Visualization Tools to Space and Atmospheric Sciences*, edited by E.P. Szuszcwicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Szuszcwicz, E.P. and J. Bredekamp, editors, *Advances in Data Handling and Visualization Tools for Applications to Space and Atmospheric Sciences*, NASA Headquarters Publication, 1993.
- Szuszcwicz, E.P., A. Mankofsky, P. Blanchard, C. Goodrich, D. McNabb, and D. Kamins, SAVS: A Space and Atmospheric Visualization Science System, in *Developments and Applications of Data Handling and Visualization Tools to Space and Atmospheric Sciences*, edited by E.P. Szuszcwicz and J. Bredekamp, NASA Headquarters Publication, 1994.
- Twiddy, B., J. Gleason, M. Schoeberl, R. McPeters, Visualization of a 3-Dimensional Ozone Data Set: Nimbus-7 SBUV (abstract), *EOS, Trans., AGU*, 74 (16), Spring Meeting Suppl., 268, 1993.

Data Handling and Visualization for NASA's Science Programs

JOSEPH H. BREDEKAMP

Office of Space Science

NASA Headquarters, Washington, DC

Advanced information systems capabilities are essential to conducting NASA's scientific research mission. Access to these capabilities is no longer a luxury for a select few within the science community, but rather an absolute necessity for carrying out scientific investigations. The dependence on high performance computing and networking, as well as ready and expedient access to science data, metadata, and analysis tools is the fundamental underpinning for the entire research endeavor. At the same time, advances in the whole range of information technologies continues on an almost explosive growth path, reaching beyond the research community to affect the population as a whole. Capitalizing on and exploiting these advances are critical to the continued success of space science investigations. NASA must remain abreast of developments in the field and strike an appropriate balance between being a smart buyer and a direct investor in the technology which serves its unique requirements. Another key theme deals with the need for the space and computer science communities to collaborate as partners to more fully realize the potential of information technology in the space science research environment.

1. INTRODUCTION

The 90s and beyond promise to be a very exciting and challenging era for conducting science from space. During this era we will be afforded the opportunity to observe the universe in all wavelengths of the electromagnetic spectrum, and compare those observations with theoretical models and simulations of the origin and evolution of the universe. We will continue to explore all the bodies in our own solar system while also seeking the existence of others. The generation of solar energy, its transport, and its interaction with the terrestrial environment will be measured. Sustained experiments in the microgravity environment of space will be conducted. And, finally, we will study the Earth as a system while determining the extent, causes, and consequences of global climate change.

We are challenged to carry out this grand, scientific endeavor within tight budget constraints, and continue to fulfill our responsibility to contribute to America's overall economic vitality. Furthermore, we are challenged to strengthen the relevance of our science mission with a more direct connection to the well-being of all Americans. It is imperative to broaden our outreach efforts to share the excitement of scientific discovery with the general public and to infuse the products of increased knowledge and understanding into educational curricula as a true legacy to the future.

Advanced data and computing systems and the full suite of related information technologies are critical, if not enabling, in achieving our space science objectives. High performance computing and networking, along with sophisticated software tools and techniques, to manipulate, analyze, and visualize complex data sets, are an integral part of the research environment for extending scientific

understanding and insight. NASA must remain abreast of this rapidly evolving area of technology to exploit advances and ensure a close coupling to the science research environment. NASA's approach to ensuring this coupling is based on three principal elements:

- The provision for state-of-the-art information systems services with reliable and expedient network access to data, easily transcended hierarchy of computing capabilities, and interactive analysis tools.
- The application of new technology to support growing demands and enable new research methods.
- The involvement of the science community from the onset to foster partnerships among computer scientists, technologists, and space scientists.

We will now go into detail about some of the challenges associated with the current and future environment, the opportunities afforded through information technology, and examples of how the above strategy is being pursued.

Figure 1 serves as a frame of reference for this discussion by providing an overview of functions and services associated with science information systems.

2. TRENDS AND CHALLENGES

Trends for future missions point toward more complex instruments with sharply increased data rates. In contrast with earlier single purpose investigations, involving observations from one instrument, future investigations will be much more interrelated in terms of dependence on multiple

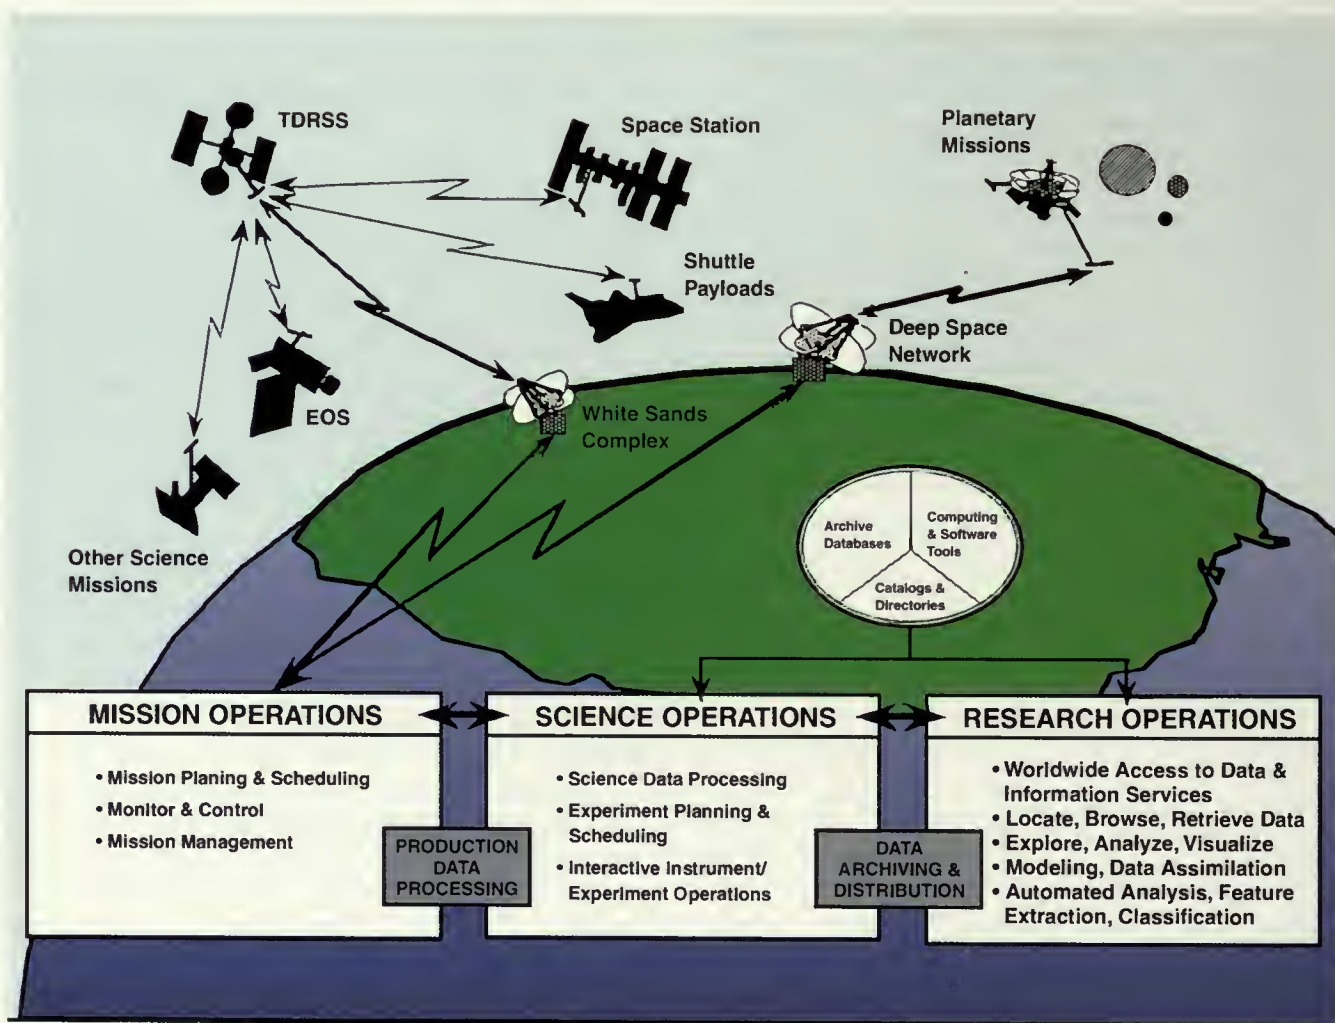


Figure 1. Functional overview of the science information systems environment.

sources of data across missions, instruments, and ground sources. Complexity of observations will also increase with investigations requiring coordinated and often simultaneous observations in multiple wavelengths with high spatial, temporal, and spectral resolution.

New patterns and research styles are anticipated for analyzing data and producing scientific findings. The value of the data assets acquired from space extends far beyond the life of the mission itself with data used in scientific studies not originally envisioned and dynamically defined as the heuristic process evolves. Collaborations are increasingly international in scope. Broad, scientific questions will be multidisciplinary in nature and involve widely dispersed investigator teams. The teams will require the combination and analysis of data from multiple sources. The need for comprehensive data archiving and distribution systems will be increasingly important. It will transcend individual project and science discipline boundaries.

The principal product of the NASA science program is increased knowledge and understanding, which are frequently represented in the form of models, which simulate and predict the phenomena and/or physical processes in question. The ultimate test of that understanding is the closure between the theory and observation, or how the process as modeled compares with observed events. Therein lies another set of challenges to the information systems environment. In striving toward the goal to represent and predict physical processes on a broad range of spatial and temporal scales, there are several factors that drive increases in computational intensity and complexity:

Resolution—Refining grid sizes to more accurately account for fine-scale effects, etc.

Sophistication—Allowing for better physics with more realistic descriptions replacing simplifying assumptions, with improved algorithms and techniques, with better parameterizations, etc.

Duration—Conducting many simulations to carry on true scientific experiments, as opposed to anecdotal studies.

Comprehension—Visualizing and interacting with results. The ability to view the event from the inside and from all aspects and

perspectives. Just as importantly, the ability to convey the results so that colleagues and laymen alike can grasp the meaning.

The combined effect of the trends and requirements in multidisciplinary space science research is virtually unique in its demands for computing power, data handling capability, and scientific visualization.

3. OPPORTUNITIES

Information technology is one of the most rapidly advancing technology areas, with profound effect on, not only research and engineering, but also on almost every household. These technologies are part of the information revolution which is stimulating economic growth and improving the quality of life. The National Information Infrastructure or “information highway” has become a part of the universal parlance and is increasingly evident in the popular news media and commercial offerings.

Trends in technical capabilities in this area are dramatic and far-reaching. Today’s workstation is yesterday’s supercomputer, which was only available from a central computing complex. Wide-area networking has revolutionized the way science is done by putting computing resources, data resources, and collaborators into a virtual laboratory setting. Similarly, very high volume data storage and management capabilities, along with affordable multi-media presentation options, are available to virtually every researcher, if not every home.

Every aspect of NASA’s science activity profits from continuing advances in information technology, ranging from operations aspects, through science data product generation, to analysis and interpretation, comparison with theory, and presentation and defense of results. In the area of on-board processing for instance, microelectronic miniaturization and packaging have opened a wide range of possibilities. Advanced flight computers with greatly enhanced processor speed, memory capacity, and reconfiguration alternatives will allow more sophisticated on-board processing for autonomous operations, intelligent image registration and classification, compression, selective downlink, etc. This technology now makes feasible flight system vs. ground system trade-offs to influence design

decisions with high potential for reducing burdensome operations costs, which is a prime target area for creating funding flexibility for future flight opportunities.

In addition, continued improvements in automation, expert systems, and telepresence methods will increase investigator interaction and remote control. These capabilities will extend to school children and the general public so they can share the experience of being there, whether to fly by a nearby planet, to get the daily weather report from Mars, or to drive a rover.

There is not a more striking example of where the integrated effect of advances in information technology combine to benefit the science process than in the area of visualization. Advanced computing engines to drive real-time rendering for three-dimensional animations; high resolution image processing and display devices; the ability to locate, retrieve, fuse, and display widely distributed data sets; and advanced algorithms and interactive analysis techniques have unlimited potential for unlocking the scientific puzzles held captive within the massive amount of data. This applies to both observational data and that generated as output by highly sophisticated models and analysis systems. It is important to note that these visualizations are more than pretty pictures. They are keys to the discovery process leading to insight and understanding and have the ability to convey that insight to others. Prospects in the area of virtual reality are also exciting with interactive exploration and experience, with the viewer embedded within the visualizations as an actual participant.

Several principles will guide the evolution of information systems support for the science research community. A portion of direct science program funding will be dedicated to assure reliable and robust data, computing, and communications infrastructure as a vital supporting element for the science mission. This will extend and build on today's NASA Science Internet to connect NASA assets consisting of computing resources, data centers, analysis tools, and science users into a worldwide network of collaborators, resources, and services.

In order to effectively apply advances in technology, a coherent and organized process needs to be employed. This process must identify, assess,

select, evaluate, demonstrate, and infuse developing technologies. Furthermore, the science community, as recipient and beneficiary, must be involved in the process from the onset. The assurance of a close coupling between information technology and science needs is predicated on team efforts involving technologists and system providers, coupled with a strong science application wrap-around. The Applied Information Systems Research component of the Science Information Systems Program was established on that basis. Figure 2 illustrates some examples of complementary research and tool developments sponsored under the auspices of that program and applied to actual science investigations. They are also representative of the underlying motivation for organizing the special session at the meeting of the American Geophysical Union, which in turn led to the publication of this book as a collective set of contributions to that session.

Another means for exploiting the technology curve for science benefit is through coordination and participation in related efforts within NASA, other federal agencies, industry, and academia. This takes the form of collaborative testbeds and technology evaluations, as well as dual-use partnerships with industry to infuse new technology into the science environment and to encourage its subsequent transfer for broadened commercial application.

An excellent example of an effective push and pull partnership between NASA science and technology programs is the High Performance Computing and Communications (HPCC) Program. This is a major national initiative, involving ten Federal agencies, to extend U.S. leadership in high performance computing technologies and accelerate the application of these technologies in the Federal government and throughout the economy. It is also intended to provide critical technology and application components for the new National Information Infrastructure. NASA's participation in this program is organized around two principal Grand Challenge applications, computational aerosciences and Earth and space science. The grand challenge in Earth and space science is to demonstrate the potential afforded by teraflop system performance to further the understanding and ability to predict the dynamic interaction of physical, chemical, and biological processes affecting the solar-terrestrial

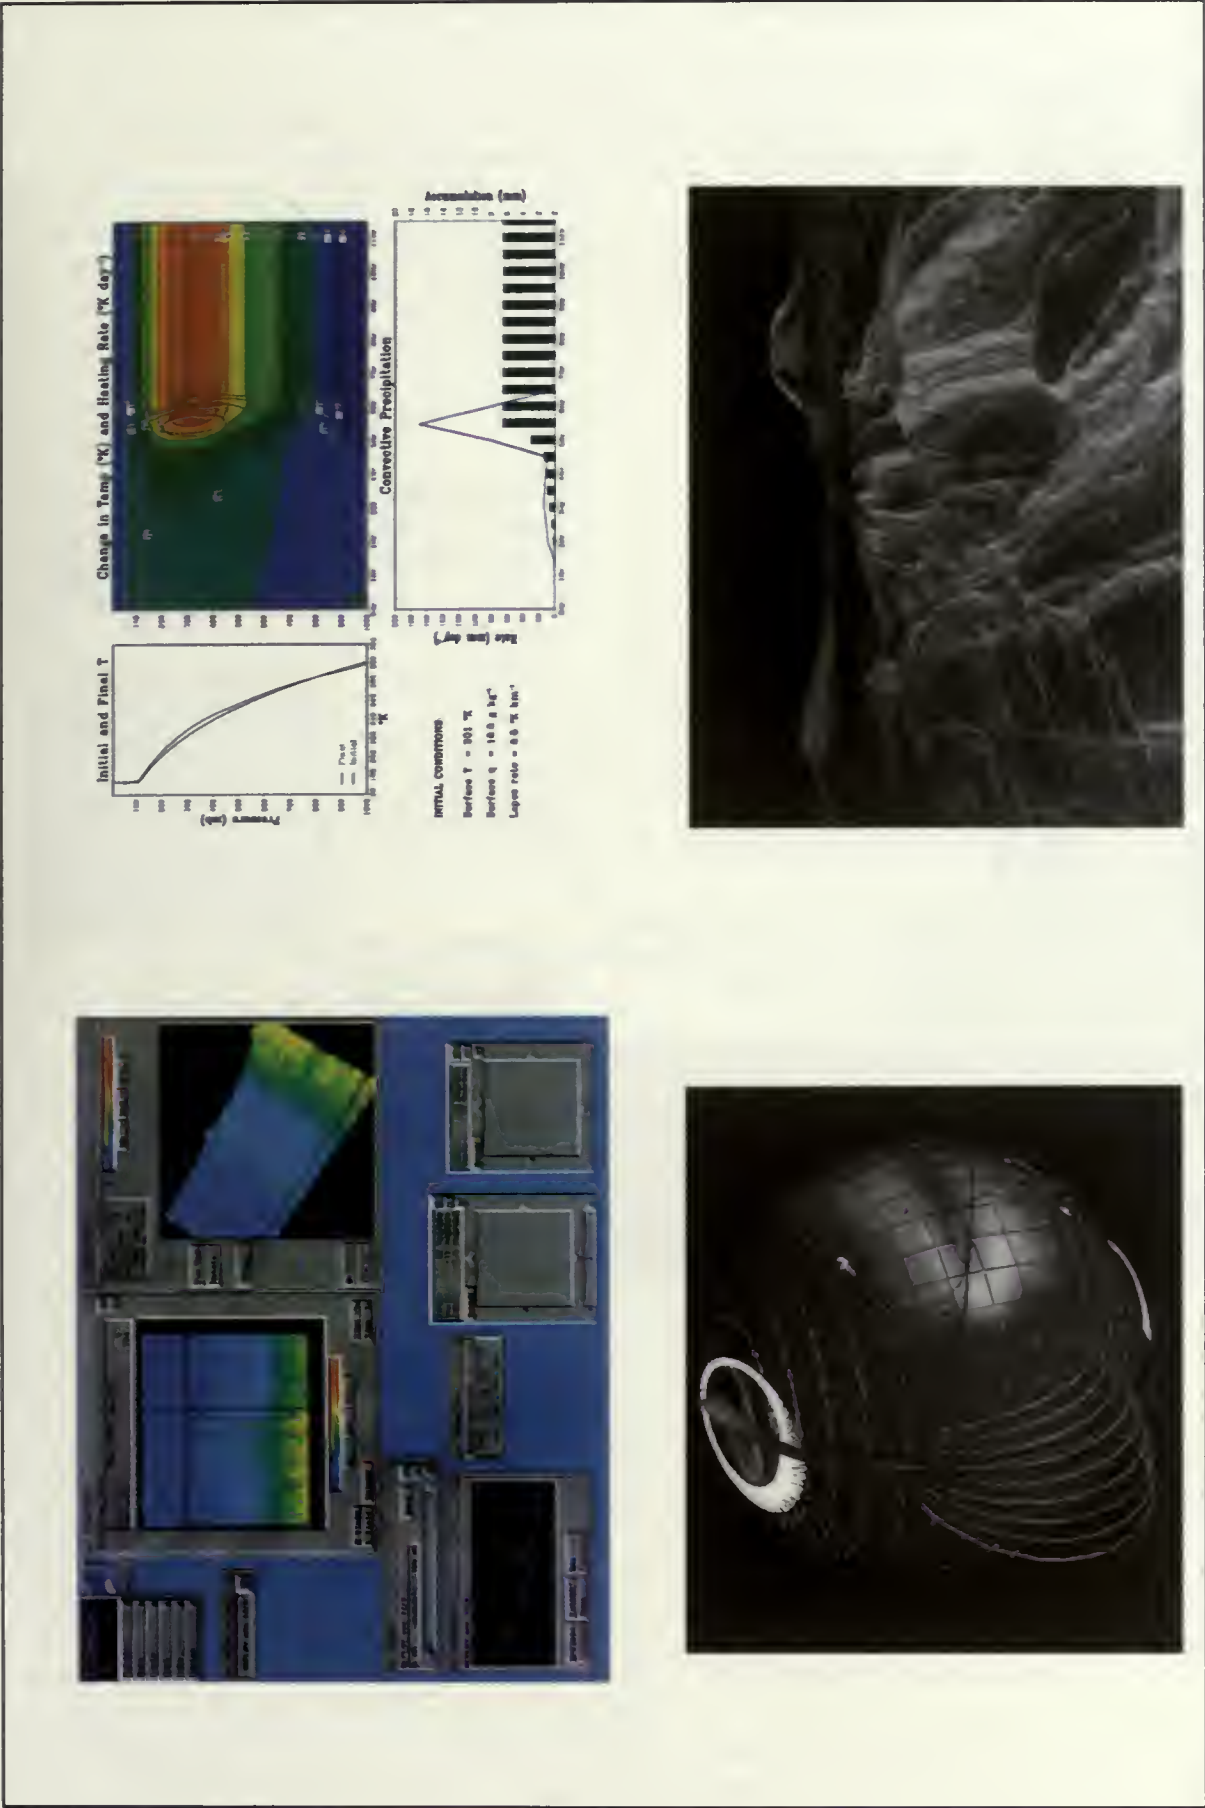


Figure 2. Illustrations using Office of Space Science-sponsored complementary research and tool developments. Clockwise, from top left: Electric field data from the Plasma Wave Spectrometer instrument onboard the Galileo spacecraft during Earth encounter, using the LinkWinds graphical spreadsheet system; Adjustment of temperature in an unstable atmospheric column by cumulus convection, as displayed using Gridded Analysis and Display System; three-dimensional rendering of the surface of Venus derived from Magellan radar imaging data (DeJong, 1994); and ionospheric-thermospheric system in support of the NASA CRRES mission, using the Space and Atmospheric Visualization Science system.

8 Grand Challenge Principal Investigator Teams

*selected through the
NASA Research Announcement
(NRA) Process*




Figure 3. Earth and Space Science Grand Challenge Applications selected through the NASA Research Announcement process for the NASA High Performance Computing and Communications Program.

environment and the universe. Twenty-nine investigations were selected for participation in the program through an open solicitation inviting proposals in two categories. Eight Principal Investigator teams of space and computer scientists were selected to adapt computer intensive models and simulations to run on experimental testbeds of massively parallel computing systems. These investigations are illustrated in Figure 3. There are also 21 smaller scale guest computational investiga-

tions focused on specific algorithms and computational methods for parallel computing architectures.

4. SUMMARY

Information systems and related technologies are essential to the furtherance of our space science research objectives. Therefore, it is critical to remain abreast of the rapidly advancing technology in order to be an early and innovative applier. It is also

important to maintain cognizance in the information technology arena in order to influence future technology developments as driven by special and/or unique needs relevant to space science.

The key to coupling the demanding space science requirements for computing power, data handling capability, and scientific visualization with the opportunities afforded through information technology is a strong strategic partnership involving the space science and computer science and technology communities. These multidiscipline collaborations build on mutual respect for the participating professions and their co-dependence for success.

Information technology is truly an explosive field. It is advancing at such a rate that it accelerates the actual realization of our potential future far more quickly than the time horizon used for the projection. And, by tradition, the science research community has pushed the envelope in exploiting advances. The combined effect offers an extremely exciting realm of possibilities.

The remaining chapters in this book provide an initial set of successful collaborations to enhance the scientific research environment, and a glimpse into the prospects for more to come.

REFERENCES

- DeJong, E.M., "Solar System Visualization: Global Science Maps," *Frontiers in Data Visualization*, London Academic Press, 1994.
- EOS Data and Information System. NASA, 1992.
- High Performance Computing and Communications: Toward a National Information Infrastructure, A Report by the Committee on Physical, Mathematical, and Engineering Sciences, Federal Coordinating Council for Science, Engineering, and Technology, Office of Science and Technology Policy, 1994.
- NASA High Performance Computing and Communications Program FY94 Program Plan, NASA, 1994.
- Report of the Information Systems Strategic Planning Project, NASA, 1990.
- Space Science for the 21st Century, NASA Office of Space Science Strategic Plan for 1995-2000, NASA, 1994.

Chapter II

**SPACE AND
UPPER
ATMOSPHERIC
SCIENCE
APPLICATIONS**

Magnetic Field Experiment Data Analysis System

**D.B. HOLLAND, L.J. ZANETTI, L.L. SUTHER,
T.A. POTE MRA, B.J. ANDERSON**

*The Johns Hopkins University Applied Physics
Laboratory, Laurel, MD*

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) Magnetic Field Experiment Data Analysis System (MFEDAS) has been developed to process and analyze satellite magnetic field experiment data from the TRIAD, MAGSAT, AMPTE/CCE, Viking, Polar BEAR, DMSP, HILAT, UARS and Freja satellites. The MFEDAS provides extensive data management and analysis capabilities. The system is based on standard data structures and a standard user interface. The MFEDAS has two major elements: 1) a set of satellite unique telemetry processing programs for uniform and rapid conversion of the raw data to a standard format and 2) the program Magplot which has file handling, data analysis and data display sections. This system is an example of software reuse, allowing new data sets and software extensions to be added in a cost effective and timely manner. Future additions to the system will include the addition of standard format file import routines, modification of the display routines to use a commercial graphics package based on X-Window protocols and a generic utility for telemetry data access and conversion.

1. INTRODUCTION

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) Magnetic Field Experiment Data Analysis System (MFEDAS) was created to provide magnetic field data access and analysis capabilities for satellite magnetic field experiment data from TRIAD, MAGSAT, AMPTE/CCE, Viking, Polar BEAR, DMSP, HILAT, UARS and Freja. The MFEDAS provides visual access to these data which total more than 50 gigabytes in size. Figure 1 is an artist's conception of the field aligned or Birkeland current system which connects the Earth's ionosphere in the northern and southern auroral zones to the magnetosphere and solar wind. The figure shows the cone-shaped model of Birkeland currents with the accompanying horizontal electrojet current system. These current systems were discovered by APL spacecraft in 1963 and subsequently studied by using magnetic field disturbances measured by spacecraft-borne magnetometers. The MFEDAS has been developed to efficiently extract information on these current systems from the instrument telemetry stream, perform the transformations needed for analysis, remove any background signals or noise that mask the features of interest, and display the analyzed data. The MFEDAS system combines and automates these tasks for efficient, individual event analysis and for rapid, automatic large scale processing.

The evolution of MFEDAS has been driven by two factors. The first is the increased rate of sampling frequency and overall data volume (TRIAD was 3 axis, 13 bits per sample and 1/2 Hz sampling, Freja is 7 channels, 16 bits per sample, 128 or 256 Hz sampling). The second factor is the need to




Figure 1. An artist's conception of the ionospheric current systems. The graph shows an example of spacecraft data measurements produced by these systems.

efficiently analyze data from a large number of satellites. During the design of MFEDAS, specific data set and application software commonalities were determined. These common elements have been incorporated into standard data structures, allowing the construction of modular, expandable software tools. Because of these standards, existing software becomes the basis for extensions of the processing and analysis software set and for the addition of the software required for new data sets. This software reuse allows these extensions to be rapidly completed with a minimum of redundant effort (e.g., FREJA required 3 staff months to incorporate into MFEDAS).

The first section below describes the logical steps in processing and analyzing these data sets. The next section describes the software required. The final section gives some examples of applications of the MFEDAS.

2. MAGNETIC FIELD DATA PROCESSING AND ANALYSIS

The MFEDAS analysis of magnetic field data is a process of several logical steps. The raw instrument data (in units of instrument counts) is converted to real physical units (typically nanoTeslas) by applying calibration and offset corrections. A time is associated with each magnetic field measurement. Erroneous data points are detected and data error handling is performed. Early quality evaluation and data screening are performed to reduce the amount of redundant error checking that would otherwise be needed in later routines. The resolution of the data may be reduced as part of the initial access to reduce the volume of the data to be analyzed. The measured data is rotated to an orthogonal XYZ instrument coordinate system and converted from instrument to science units. The

values for the conversion factors are estimated by measurements on the ground and refined by using measurements during the mission. Once the data are in science units, they are rotated into the standard spacecraft referenced coordinate system. The final data transformation is made from the spacecraft reference system into an Earth or star-fixed reference system by applying transformations that incorporate spacecraft attitude and location data.

Analysis often includes the tasks of separating the ionospheric induced fields from the measured field which is a composite of the main geomagnetic field, the ionospheric induced fields, and other noise sources within the instrument and spacecraft. The first step in separating the field disturbances from the measured values is the subtraction of a model of the geomagnetic field. The spacecraft position and attitude must both be well known to do this, because the disturbances are only a small part of the measured field on the order of a few percent of the Earth's field. The data are transformed into common coordinates and common displays to allow the data to be compared with other data sets and to best display specific features.

3. DATA PROCESSING

The MFEDAS provides the software tools needed for magnetic field experiment data set processing and analysis tasks described in the preceding section. Figure 2 shows the two major processing steps of the system and data flows through the system. The first step contains satellite unique telemetry processing software designed to allow users to read and summarize the resulting data and to obtain uniform and rapid data conversion to a standard format. The second major processing step, performed by the program Magplot, includes the set of general routines to perform the tasks required for data analysis and display for all data sets. Magplot includes common file access and file manipulation routines, analysis software including filtering, averages, fits, noise reduction and FFT's and displays. The display options include multiple data item plots, polar plots and spectrograms. Magplot has internal standards for magnetic field data structures, ephemeris data access, and for calling parameters. These standards are the basis for reuse of the existing software for new data sets and for extensions to the software set.

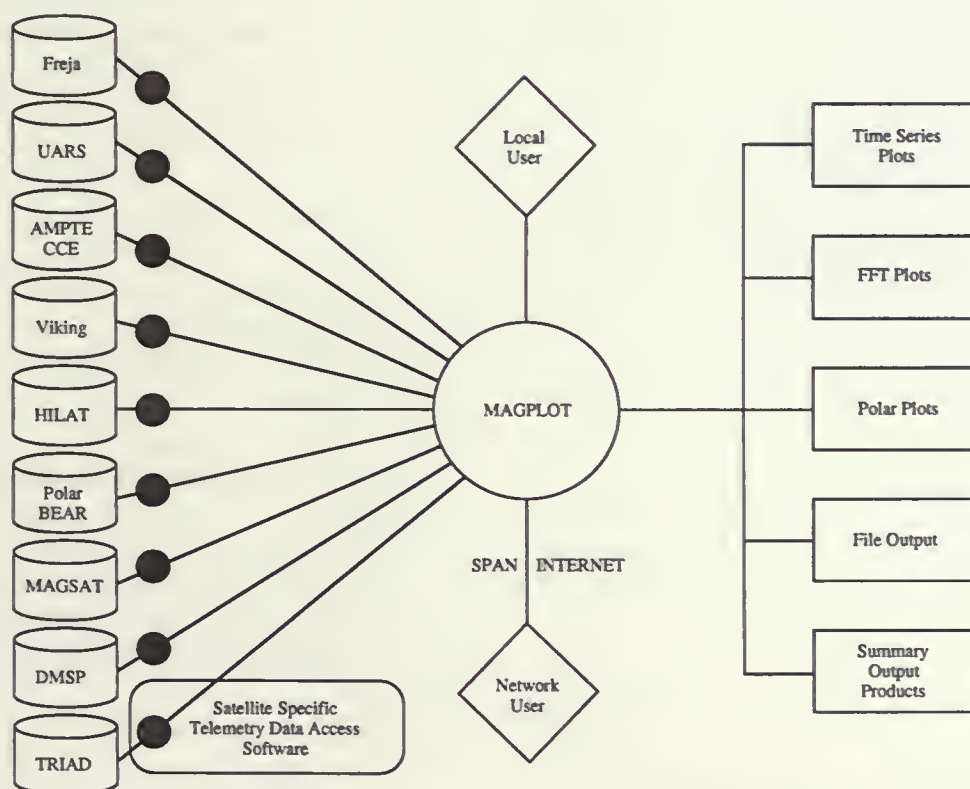


Figure 2. The figure shows input data sets, processing modules and products of the MFEDAS. On the left are the satellite unique processing software modules. The central part of the MFEDAS is the generic software in the program Magplot, which can be reused for new data sets.

Each new data set requires a unique software data extraction module to be created to deal with the variations in the input data sets. Satellite data sets are made available on a range of physical media and in unique satellite telemetry stream data formats. The extraction software contains the locations and data types of the required fields in the input data stream and the operations required to convert these fields to a usable value. These operations include bit, byte and word access and conversion to local types, real data representation conversions, and other data type conversions. The large input data stream size requires that these operations be done in an efficient manner. Developing this software can be the most time consuming part of providing access to a new data set.

A planned extension to the MFEDAS is a generalized utility that uses a graphic interface to interactively create the mapping between the data stream and standard format output files. This utility would make available an extensive set of tools for these operations. The utility would remove the inefficiencies of individual software creation by providing a rapid and self documenting method for low level data access.

Figure 3 diagrams the layers of the logical shell structure of the general analysis and display program Magplot. Magplot consists of a core data structure and operators on this structure. The three segments of operations share the core data structure. The operations are accessed through an interactive user interface or the command script interface. The labels in Figure 3 are described in detail in the paragraphs below.

3.1 Standard Data Structure and Data Structure Handling

All routines use a common data structure. This structure consists of a descriptive header and a multiple record data area. The header consists of 24 integer descriptive values that encode the type of data, satellite that creates the data, coordinate system, times and further descriptive information. The data area consists of a column of time (seconds of the day) and up to nine columns of data. A set of routines has been developed to display and modify the header and to display, scale, and reorder the data structure.




Figure 3. Shell structure diagram showing the logical layers in the program Magplot.

3.2 Data File Handling and Ephemeris Data Access

Routines are provided for input and output of standard format data files. Ephemeris data is available from within Magplot by standardized calls that input the data from disk files. A data base of the ephemeris data available is used to provide rapid, automatic file selection.

3.3 Data Analysis and Filtering

Extensive analysis software is available including filtering, averaging, fits, noise reduction, FFT's, plots and spectrograms multiple filters, polynomial fits, and averages.

3.4 Graphic Creation and Display Handling

Displays include multipanel time series plots, polar dial plots, and spectrograms. Output display device drivers are available for X-Windows, TEK4010, QMS Laser printer, and PostScript.

3.5 User Interactive Interface

The user interface is flexible, menu driven, and rapidly modifiable. The tree structure for command

selection allows maximum user flexibility. Figure 4 shows the first two levels of menus.

3.6 Command Script Interface

Prototype command scripts have been set up to allow rapid creation of scripts for specific tasks. These are used to create specific command scripts to be used for production processing.

The MFEDAS includes a set of modules to perform coordinate system transformations for standard coordinate systems: geographic and geomagnetic; earth centered and local; and Earth, solar and stellar fixed.

4. EXAMPLES OF THE ANALYSIS OF MFEDAS RESULTS

Figure 5 is a composite plot of six of the JHU/APL spacecraft data sets produced by the MFEDAS. The plots are polar plots in magnetic latitude versus magnetic local time of disturbance vectors transverse to the main field with bases along the orbit trajectory. The coordinate systems shown have one component parallel to the magnetic

field and the other two orthogonal components are in the North and East or Sun and Dawn directions.

The UARS magnetic field data of the upper left panel shows transverse disturbance vectors indicating strong Region 1 (middle cone of Figure 1) and Region 2 (equatorward cone of Figure 1) Birkeland current systems associated with an expanded auroral oval ($\sim 60^\circ$ MLAT circle). These large scale dynamics are common for the present solar cycle maximum. The third line trace of the UARS panel is a plot of magnetic disturbances parallel to the main field indicating the ionospheric electrojet currents drawn schematically as extended horizontal arrows. The UARS panel is a composite of northern and southern hemisphere passes.

The FREJA panel shows a sample of recent data showing the dayside Region 1 and Region 2 Birkeland current disturbances during an active magnetic storm development on December 28, 1992. Polar cap disturbances have been artificially nulled.

The MAGSAT panel is a composite of three polar passes during northward IMF with the transverse vectors indicating the polar cap NBZ Birkeland current system (inner cone of Figure 1),

<u>MAGPLOT</u>		
<u>MAGFIL</u> File access routines	<u>MAGPRO</u> Data processing routines	<u>MAGPLOTS</u> Plotting routines
MAGREAD: Input standard format magnetometer data	GEOGM: Rotate from satellite to geographic coordinates	XYPLOT: Time vs. B(1-4) or Frequency vs. Amplitude
MAGWRIT: Write a file of the current data	FIX: Remove bad data points with automatic selection	POLARPLOT: Polar plot & three panel T vs. B
READDAT: Free format file input	SLAV: Sliding average	MLATB: Magnetic latitude vs. B
BLST: List header variables	AVG: Average	POL6: Polar vector plots
CMNT: Change header variables	TAPR: Taper data	
SAMP: Create sample data set	DTRN: Detrend data	
SWAP: Move main data arrays to/ from secondary arrays	POLFIT: Detrend data using a polynomial fit	
OAOPS: Call att. data access routines	FILT: Bandpass or bandstop filter using Fourier transforms	
OATAG: Add att. data to data	FRTR: Fourier transform	
OALIST: List all att. data file	IFTR: Inverse Fourier transform	
OAINDEX: Access att. data index	TRNC: Limit time range of data	
GETOA: List att. data	POWR: Spectral power density	

Figure 4. Command menus and command names available in Magplot. The First level shows the three program divisions. The second level shows command options.

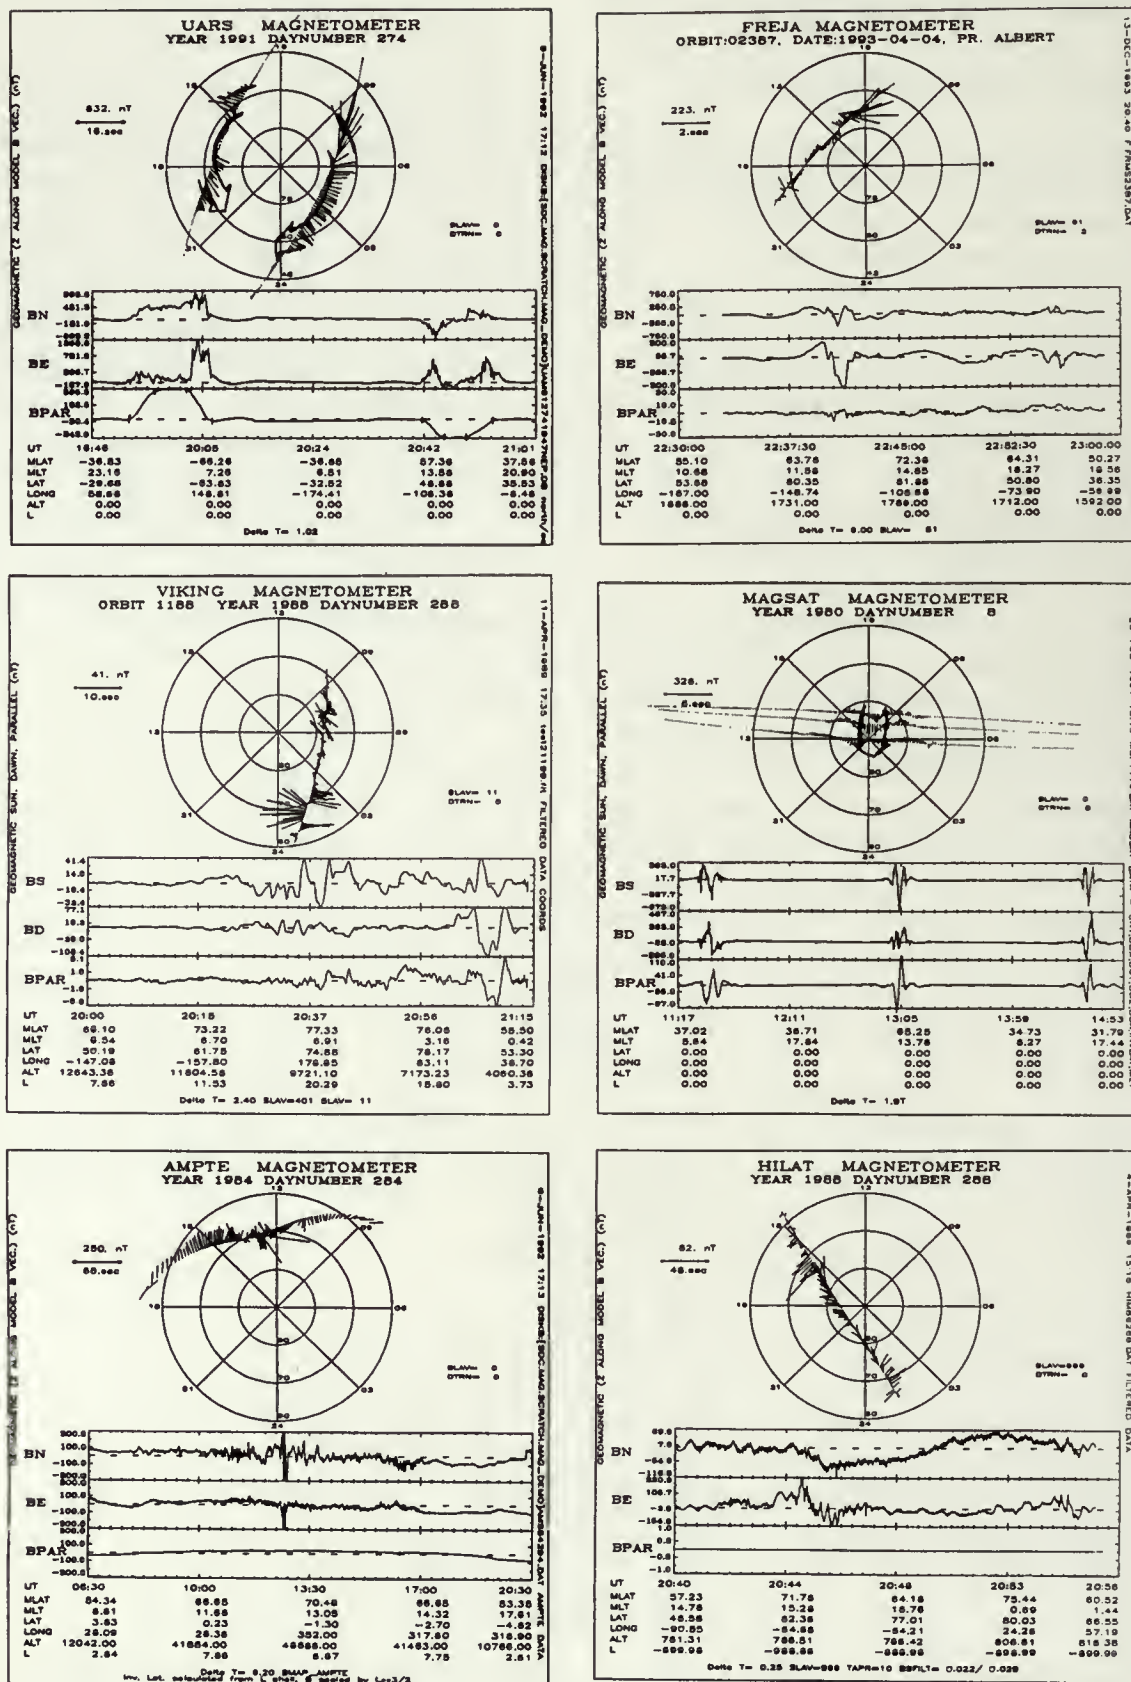


Figure 5. A composite figure showing plots of six data sets processed with the MFEDAS.

which has been proposed to be driven from neutral wind inertia and thermosphere-ionosphere momentum, coupling back to the ionosphere during quiet times following magnetic storm activity. The magnetic field disturbances parallel to the main field indicate the antisunward ionospheric Hall current associated with these events and are indicated schematically.

The Viking and HILAT panels are correlative data from the two spacecraft at the same time. Transverse disturbance vectors are displayed which infer the large scale Birkeland current system during this transition from northward IMF into the storm time expansion phase.

The AMPTE panel shows data mapped along magnetic field lines from the CCE equatorial orbit on the dayside indicating the ionospheric location of magnetosheath/cusp field disturbance signatures.

Figure 6 shows polar plots of six hours of magnetic field data created with a MFEDAS extension that uses the information in the AC channel parameter file to automatically produce plots during periods of high activity. The NASA Upper Atmospheric Research Satellite (UARS) was launched September 12, 1991, with an APL magnetic field experiment. The unique stability of this platform and the full orbit data collection have provided geomagnetic observations from UARS that approach MAGSAT quality. The standard production of the UARS magnetometer summary data includes the processing of the 5 - 50 Hz peak detector channel. This circuit responds to wave phenomena in that frequency range and because of the wave turbulence attendant to the Birkeland currents and the high frequency power in the Birkeland current signature discontinuities. This




Figure 6. Polar plots of six hours of magnetic field data from the UARS spacecraft showing magnetic field activity during successive orbits on April 5, 1993. The time periods for these plots were selected automatically from the activity index values and the command script interface to Magplot was used to control file access and plotting.

AC channel serves as an indicator of spacecraft passage into the auroral zone, reflecting both the intensity of the currents and their duration. During automatic production processing, a full day plot is produced and a condensed on-line ASCII file of AC channel parameters is created, which gives start and stop times and locations of auroral zone crossings above a significant threshold level. The FREJA magnetic field experiment has a similar "event detector" whose data can be received real-time or transferred via network from Sweden.

5. SUMMARY

The MFEDAS is a comprehensive system for efficiently accessing and analyzing magnetic field data. The system is an example of software reuse where the sophistication and breadth of the tasks that the software performs have grown with the number of data sets available to the system. This system growth is achieved by maintaining standard data structures and interfaces. This has allowed the accumulation of knowledge about the requirements of magnetic field data analysis to be incorporated into the MFEDAS software.

The MFEDAS has successfully integrated the nine scientific magnetic field data sets into a common analysis system. Although the initial development of this system took a typical amount of resources and level of effort, subsequent data sets, such as UARS and Freja, have been integrated for a small fraction of a staff-year and using minimal resources. The built-in versatility has proved invaluable to troubleshooting data, spacecraft problems, and variations inevitable with each new data set. The MFEDAS is a permanent resource, providing new projects with a quick jump start to data analysis and science investigations by decreasing software development time to a minimum and reducing costs to sponsors.

REFERENCES

- Anderson, B.J., T.A. Potemra, P.F. Bythrow, L.J. Zanetti, D.B. Holland, and J.D. Winningham, Auroral currents during the magnetic storm of November 8 and 9, 1991: observations from the Upper Atmosphere Research Satellite Particle Environment Monitor, *Geophys. Res. Lett.*, 20, 1327, 1993.
- Potemra, T.A., Field-aligned (Birkeland) currents, *Space Science Reviews*, 42, 295, 1985.
- Potemra, T.A., L.J. Zanetti, and M.H. Acuña, The AMPTE CCE Magnetic Field Experiment, *IEEE Transactions on Geoscience and Remote Sensing*, GE-23, 246, 1985.
- Potemra, T.A., L.J. Zanetti, R.E. Erlandson, P.F. Bythrow, G. Gustafsson, M.H. Acuña, and R. Lundin, Observations of large scale Birkeland currents with Viking, *Geophys. Res. Lett.*, 14, 419, 1987.
- Potemra, T.A., Observation of Birkeland currents with the TRIAD satellite, *Astrophys. and Space Sci.*, 58, 207-226, 1978.
- Zanetti, L., T. Potemra, R. Erlandson, P. Bythrow, B. Anderson, A. Lui, S. Ohtani, G. Fountain, R. Henshaw, B. Ballard, D. Lohr, J. Hayes, D. Holland, M. Acuña, D. Fairfield, J. Slavin, W. Baumjohann, M. Engebretson, K. Glassmeier, G. Gustafsson, T. Iijima, H. Luehr, F. Primdahl, Magnetic Field Experiment for the Freja satellite, *Space Sciences Reviews*, (in press), 1993.
- Zanetti, L.J., W. Baumjohann, and T.A. Potemra, Ionospheric and Birkeland current distributions inferred from the MAGSAT magnetometer data, *J. Geophys. Res.*, 88, 4875, 1983.

Scientific Visualization Tools for the ISTP Project: Mission Planning, Data Analysis and Model Interpretation

M. PEREDO

*Hughes STX Corporation at NASA Goddard Space Flight
Center, Code 695, Greenbelt, MD*

C.C. GOODRICH, D. McNABB, R. KULKARNI

*Department of Astronomy, University of Maryland,
College Park, MD*

J. LYON

*Department of Physics and Astronomy, Dartmouth
College, Hanover, NH*

Visualization tools are being developed to meet the challenges of mission planning and data analysis presented by the International Solar-Terrestrial Physics (ISTP) program. ISTP encompasses a large number of spacecraft, multiple ground-based observatories and several theoretical investigations, with the goal of understanding the global behavior of the solar wind-magnetosphere-ionosphere system. The tools include three-dimensional displays of key boundaries in geospace along with spacecraft trajectories, which can be animated and synchronized to universal time. Magnetic field models and MHD simulation results can be invoked to reveal the magnetic topology or to identify magnetic conjunctions between spacecraft and/or ground-based facilities. Simultaneous displays of satellite trajectories, spacecraft-borne observations, and model predictions are available to facilitate data processing and interpretation efforts. The current status of these tools is described, and their implementation at the ISTP Science Planning and Operations Facility and distribution to the entire ISTP community are discussed.

1. INTRODUCTION

The International Solar-Terrestrial Physics (ISTP) program seeks to understand the physical processes that transport mass, momentum and energy from the Sun, through the interplanetary medium, into and through the Earth's magnetosphere, and finally into the ionosphere. Specifically, the program focuses on the study of plasmas and magnetic fields of the Sun, the Earth, and the space in between. The program involves simultaneous, coordinated measurements from key areas of geospace. The scope and magnitude of the ISTP mission demand a sophisticated set of visualization tools to perform operational planning and collaborative analysis of data from space-borne instruments, ground-based observations, and empirical or theoretical models.

Visualization of ISTP data is particularly complicated in two respects. First and foremost is the inherently global nature of the project, which demands integrated views of the data. Scientists are faced with the challenge to visualize and analyze data collected by different instruments onboard multiple spacecraft situated in vastly different regions of geospace, and to integrate these data with those obtained from ground-based or ionospheric measurements. While theory and modeling efforts provide the "maps" to facilitate comparisons and calibrations of observations from disparate sources, interactive visualization tools are essential to understanding the spatial and temporal relation between the observations in the context of the entire solar wind-magnetosphere-ionosphere system. Access to effective scientific visualization tools is now possible due to the development of powerful workstations. The second factor complicating the visualization of ISTP data is the unprecedented

volume and diversity of data expected from ISTP investigations. The ISTP data pool includes observations from ISTP satellites, as well as data from other spacecraft missions which have established research collaborations with the ISTP program. In addition, there are ground-based and theory investigations. Complementary data will be integrated through coordinated science campaigns with other major international science initiatives. All together, some twenty spacecraft, tens of ground-based observatories, and several theory investigation groups will provide an estimated 2.2 gigabytes of data per day [see Mish et al., 1994 for a detailed description of the ISTP data systems and products]. In this communication, we describe the current capabilities of the Visualization Tools developed for the ISTP program, their implementation at the ISTP Science Planning and Operations Facility (SPOF) [Peredo, 1993], and enhancements currently under development.

2. THE ISTP PROGRAM

As indicated above, the ISTP program is an international, multi-spacecraft project aimed at developing quantitative and predictive models of the flow and transfer of particles, momentum, and energy from the Sun, through interplanetary space into the magnetosphere and eventually into the Earth's ionosphere. To accomplish these goals, the ISTP program involves simultaneous coordinated observations throughout key regions of geospace including the polar auroral regions, the equatorial magnetosphere at geosynchronous distance, the geomagnetic tail, the magnetosheath and the solar wind. Complementing the in-situ spacecraft measurements are remote sensing data from ground-based observatories, and model and simulation results from theoretical investigations.

Spacecraft observations of the Sun will be provided by SOHO; the solar wind state will be primarily monitored by the WIND spacecraft; and, IMP-8 will provide additional solar wind data during a portion of its orbit, and magnetosheath, and tail observations in the remainder of each orbit. The POLAR and INTERBALL-AURORA spacecraft will concentrate on the Earth's polar auroral regions, while the GEOTAIL and INTERBALL-TAIL probes will sample the geomagnetic tail. The equatorial region is being covered by existing

missions from the DOE Los Alamos National Laboratory spacecraft and the NOAA Geostationary Operational Environment Satellite program. The four-satellite CLUSTER mission will provide detailed observations of small-scale, fast phenomena in the magnetosphere, magnetosheath, and solar wind.

Ground-based observations will be contributed from the Canadian Auroral Network for the Origin of Plasmas in Earth's Neighborhood Program Unified Study (CANOPUS), by the Satellite Experiments Simultaneous with Antarctic Measurements (SESAME), the Sondrestrom Radar, and the Dual Auroral Radar Network (DARN). A wide variety of ground-based instruments are used to measure, among other things, particle precipitation and flux, magnetic-wave intensity and polarization, plasma velocity and convection, and electric fields. The ground-based observations serve a dual purpose. They provide remote sensing of the magnetospheric space separating the spacecraft, while also providing high-latitude ionospheric measurements that serve to define key boundary conditions required to couple the ionospheric and magnetospheric components of theoretical models.

An innovative aspect of the ISTP program is the role played by theory and modeling investigations. The theoretical models provide a framework for understanding the spatially scattered and diverse spacecraft and ground-based observations. The observations in turn provide a standard to evaluate and improve the models. Theoretical investigations are being pursued on three levels. The global structure of the solar wind-magnetosphere system is being modeled by large scale magnetohydrodynamic (MHD) computer simulations and by empirical magnetic field models. The structures of important boundary layers, such as the magnetopause, are being studied using hybrid simulation codes and analytic models. The microscale plasma processes that may determine mass, momentum, and energy exchange across these structures are being investigated with kinetic and fluid plasma theory and computer simulation.

3. THE ISTP/SPOF

The ISTP Science Planning and Operations Facility (SPOF) [Peredo, 1993] is the component of the Global Geospace Science program responsible

for the development and coordination of ISTP science planning operations. The SPOF operates under the direction of the ISTP Project Scientist and is responsible for the development and coordination of the science plan for ISTP spacecraft. SPOF executes ephemeris and model software to generate science planning and event times, and provides the link for WIND and POLAR scientists to submit experiment operating instructions. The SPOF reviews Key Parameter data (low resolution, ~1 min., survey data), and aids investigators in the identification of intervals of scientific interest for detailed study. SPOF scientists also collaborate with ISTP Theory Investigators on the development, evaluation, assessment, and comparison of magnetospheric models. While the SPOF workstations are restricted to SPOF personnel, tools developed for mission planning and data analysis, including those described in this communication, will be made available—without any support commitment—to the entire ISTP community via the ISTP Central Data Handling Facility. The tools described in this communication will also be available via anonymous ftp from avl.umd.edu. Note that interested users must have the appropriate commercial licenses for AVS, PV-Wave and IDL in order to run the tools.

4. THE ISTP VISUALIZATION TOOLS

Understanding the critical role of visualization for ISTP, the SPOF, and the University of Maryland Advanced Visualization Laboratory [Goodrich and McNabb, 1992] are jointly working to develop interactive tools to meet the science planning and data analysis needs of ISTP [Goodrich et al., 1993, McNabb et al., 1993]. These tools are based on widely used, modular and extensible applications including the Application Visualization System (AVS) [Upton et al., 1989], the Interactive Data Language (IDL) [IDL, 1993], and PV-Wave [PV-Wave, 1993]. The modular nature of these packages means that ISTP investigators can combine efforts to build better tools. The SPOF and the AVL act as the central point for integration and distribution of the tools, but modules developed by other scientists can easily be incorporated into the tools. While the tools have been developed and tested on SUN and DEC unix workstations, the

highly portable nature of AVS, IDL, and PV-Wave should make it easy to install and run on any platform for which these products are available.

Current capabilities of the ISTP visualization tools include generation of integrated three-dimensional displays of satellite trajectories, empirical models of key magnetospheric boundaries, and magnetic field lines from empirical models or magnetohydrodynamic simulations. The spacecraft, boundaries and magnetic flux tubes may be animated, synchronized to universal time, and the ionospheric foot points may be displayed interactively. Simultaneous display of in situ, ground-based and simulation data are readily incorporated. In the following sections, we describe three AVS-based prototype applications that illustrate the current state of the ISTP visualization tools. In addition, we discuss their role for mission planning and data analysis functions at the ISTP/SPOF.

4.1 *Spacecraft and Magnetospheric Regions Display*

In order to meet the global needs of the ISTP program, our visualization tools allow the user to generate three-dimensional displays of spacecraft trajectories within the context of the entire geospace system. We have constructed three-dimensional surfaces representing boundaries between key magnetospheric regions. These boundary representations are largely based on empirical, data-based models and are compatible with those used by the NASA Satellite Situation Center (SSC) [SSC, 1992; Peredo et al., 1992a.; Parthasarathy et al., 1994; Aist-Sagara et al., 1993]. The SSC tools are used extensively at the SPOF and are complementary to the visualization and data analysis tools described in this communication; as such, use of compatible definitions for the key regions of geospace is of utmost importance. The boundary representations have been revised and updated extensively over the last two years. Our visualization tools employ an important subset of the regions defined for SSC applications; the three key boundaries included in our three-dimensional magnetospheric displays are: (1) the magnetopause model of Sibeck et al. (1991), (2) a modified version of the Fairfield (1971) bow shock model; this modified version was introduced as part of the

update of the SSC applications; the original bow shock surface derived by Fairfield is modified by allowing the surface to move along the Sun-Earth direction in response to variations in the solar wind dynamic pressure subject to the constraint that the bow shock stand-off distance is 1.3 times the magnetopause stand-off distance (see SSC, 1993 for details), and (3) the neutral sheet model of Fairfield (1980). Close interaction exists between the SPOF and SSC groups, and any future updates or revisions to the magnetospheric region definitions in the SSC and visualization tools will be coordinated. Users may replace these boundary representations with those of their choice by creating AVS geometries corresponding to the model(s) of interest and using our existing modules as example code.

Our first example AVS application produces the three-dimensional spacecraft and geospace displays discussed above. In addition to spacecraft trajectories and the magnetospheric boundaries, our displays may include orienting features, such as continental outlines on the Earth and either Geocentric Solar Magnetospheric (GSM) or Geocentric Solar Ecliptic (GSE) coordinate axes. Satellite

trajectories are readily generated from predictive or definitive ephemeris files in the Common Data Format (CDF) adopted by ISTP for many of its data products. Displays can be easily animated to illustrate the motion of the satellites and the displacement of the boundaries in response to variations in the model parameters, including solar wind dynamic pressure or in the orientation of the Earth's magnetic dipole. For these animations, all objects in the display are synchronized with respect to universal time.

Using this AVS application, we can interactively visualize and identify advantageous spacecraft configurations. This application will be used during mission planning activities to identify potential windows of opportunity to address specific science questions, and the experimenters will be alerted to configure their instruments accordingly. During data analysis efforts, this AVS network will help in the interpretation of simultaneous observations from multiple spacecraft-borne and ground-based instruments. Figure 1 shows a sample predicted configuration including the GEOTAIL, INTERBALL-TAIL, INTERBALL-




Figure 1. Sample spacecraft configuration for global studies of magnetotail dynamics. IMP-8 (green) acts as the solar wind monitor, INTERBALL-AURORA (yellow) monitors the polar auroral region, GEOTAIL (pink) samples the deep tail while INTERBALL-TAIL (dark red) explores the high-latitude magnetotail; the moon is shown (larger red sphere) for reference.

AURORA, and IMP-8 spacecraft. The image was generated during a planning session for the First Science Campaign of the Inter-Agency Consultative Group [Whipple and Lancaster, 1992], and illustrates a desirable configuration for magnetotail dynamics investigations. IMP-8 provides solar wind information, AURORA samples the polar cap regions while TAIL and GEOTAIL cover, respectively, the high-latitude and deep geotail regions.

The AVS network used to create Figure 1 contains the module we have developed to generate the surfaces corresponding to the key geospace boundaries detailed earlier; entire surfaces may be shown, or half-surfaces can be displayed where the slice is in the GSM x-z plane. Each surface is a separate geometric object, and the AVS built-in functionality to modify object properties can be used to individually assign color and transparency properties to yield the preferred view. This surface module has been enhanced to conceptually display the *windsock effect* [Hones et al., 1986] illustrated in Figure 2, and thought to occur as a change in solar wind direction upstream of the magnetosphere

reaches the magnetospheric boundary layers and propagates downtail along the magnetosphere. The location, angle, and extent of the solar wind disturbance can be controlled interactively, and it is easy to produce an animated three-dimensional display of the propagation effect. The qualitative representation in Figure 2 actually agrees well with the quantitative simulation study of the windsock effect by Goodrich and Lyon [1992].

Each of the satellites shown in Figure 1 (and their trajectories) are generated by a copy of the satellite macro module that reads the ephemeris data, automatically performs the necessary time and coordinate transformations, and produces two output fields: the spacecraft position at the selected universal time, and the list of points corresponding to a trajectory segment "about" the current position. The user can interactively control the duration of the trajectory segment. A placement slider parameter selects the portion of the trajectory segment appearing ahead of and behind the spacecraft. Another custom module allows selection of the current universal time, and permits proper synchro-




Figure 2. Simulation of the Windsock Effect. A change in direction of solar wind velocity produces a kink that propagates along the bow shock, magnetopause and neutral sheet surfaces. Our AVS application allows interactive control of the extent and sharpness of the kink to simulate specific variations in the upstream conditions.

nization of all objects in the display. A library of conversion routines has been created to facilitate conversions between different formats commonly used by the NASA community to represent time. The library also includes routines for time-arithmetic necessary for efficient data searches and selections. A macro module composed of public domain modules is used to display the Earth as a sphere with overlaid continental outlines and latitude and longitude contours. The interactive three-dimensional controls of AVS are especially helpful in this network when examining in detail such events as the crossing of one of the surfaces by a satellite. The user can easily rotate, translate, or zoom the display to view the event from multiple angles and assess its importance.

4.2 Interactive Display/Identification of Magnetic Conjunctions

The second AVS application we discuss is based around our magnetic field line tracing module. The underlying code includes routines for the

widely used Tsyanenko models [Tsyanenko 1987, 1989, 1990], as well as improved variants of these models recently derived [Peredo et al., 1992b; Peredo et al., 1993; Stern and Tsyanenko, 1992]. The module permits generation of three-dimensional displays of magnetic field lines traced from one or more seed points. Synchronization to universal time is inherently built into the network. One use of the module involves tracing lines from a large number of seed points on the surface of the Earth out into the magnetosphere; the resulting display shows the global structure of the magnetic field according to the selected model, and may be used to identify graphically the boundary between closed (both ends reach the Earth) and open (only one end reaches the Earth) field lines. An alternative use is to select a spacecraft position as the seed point, and to trace the magnetic flux tube passing through the spacecraft down to the Earth's ionosphere, as illustrated in Figure 3. Such a display, combined with an overlay of ground-based observatory locations and continental outlines on the Earth, readily provides an interactive tool for visual




Figure 3. Magnetic conjunction between the POLAR (red) spacecraft and a ground-station of the CANOPUS array. (red dots on the surface of the Earth); the SAMPEX probe (green) maps to higher latitudes. Field line traces are according to the Tsyanenko 89 magnetic field model corresponding to highest magnetic activity level ($K_p \geq 5$).

identification of magnetic conjunctions. These conjunctions, either between a spacecraft and a ground station, or between multiple spacecraft hold special scientific significance since charged particles travel preferentially along magnetic field lines, and thus magnetic conjunction times are prime candidates for effective correlative studies between ground-based and spaceborne experiments.

4.3 In Situ Data Display/Comparison

To provide researchers with powerful integrated visualization and analysis tools, we complement the three-dimensional spacecraft and boundaries displays with simultaneous time-series plots comparing predicted theory or model “data” to the actual measurements obtained by the spacecraft. As with our previously described applications, the spacecraft and boundaries motion, and the model and observed data on the time-series displays, are all synchronized to universal time.

To display the time-series data, we use AVS modules developed with the SAVS program [Szuszczewicz, et al., 1994] that call PV-Wave and use its functionality, in this case to produce stacked

line plots. The modules are linked to PV-Wave at compile time, exchange data and communicate with PV-Wave through shared variables, and directly call PV-Wave routines to execute arbitrary command language statements. Because the ISTP community is divided among their access to PV-Wave or IDL, we are working with Research Systems Inc. to develop a similar module that communicates with IDL. An illustration of our joint displays appears in Figure 4, where magnetic field components observed by IMP-8 are compared with the results of an MHD simulation. For simplicity in this example, we have interpolated the simulation results from a single time step along the IMP-8 orbit. The field components show very good general agreement, despite the limitation of actually comparing spatial variation in the simulation results with the time series of IMP-8 data. Though more arduous, we can and will be including the time dependence of the MHD results in these comparison. This application allows superposition of data and results from several sources; such data may come from empirical models such as the Tsyganenko magnetic field models [Tsyganenko, 1987, 1989, 1990], from numerical simulations,




Figure 4. Simultaneous Display of Model/Spacecraft Data and three-dimensional Magnetospheric Configuration. The three-dimensional display shows IMP-8 exiting the magnetosphere on the dawn side, while the stacked data plots show the magnetic field components observed by the spacecraft (white) and predicted by an MHD simulation (blue). These data plots were generated with PV-Wave using data imported from AVS via the PV-Wave/AVS module developed with the SAVS program [Szuszczewicz, et al., 1994].

such as the MHD simulations developed by the ISTP Theory Group [Raeder, 1992; Convery et al., 1992; Goodrich and Lyon, 1992; Lyon and Fedder, 1992] or from actual spacecraft observations.

5. SUMMARY AND FUTURE PLANS

A suite of visualization and data analysis tools is under development for the ISTP project. Several prototype applications have been created and are in active use by the ISTP/SPOF. Distribution of these tools to selected ISTP investigators with AVS experience is currently underway. Their experience with the tools and feedback will be used to determine additional features desirable to tailor the applications to the true needs of space scientists. In addition, we plan to incorporate enhancements suggested by our own use of the tools within the SPOF and ISTP Theory Groups. For example, we are developing more sophisticated CDF-readers to handle multiple ephemeris data sources and diverse key parameter data from all ISTP investigations. Future plans also include seamless integration between AVS, IDL, and PV-Wave to allow data exchange between any of these products already widely used by the space physics community. Finally, we plan to expand our currently limited distribution of the tools to include all members of the ISTP community who request these tools. Distribution of the tools is currently on an "as-is" basis, with minimal documentation and no support commitment.

Acknowledgments. We gratefully acknowledge contributions by C. Mesztenyi in the development of the boundaries module. We thank N. A. Tsyganenko for the subroutines for his models, the Satellite Situation Center for ephemeris data, and the ISTP Central Data Handling Facility for Key Parameter data. Finally, we thank Advance Visual Systems, Inc., Research Systems, Inc., the PV-Wave group at Visual Numerics, Inc., and Digital Equipment Corp. for technical assistance.

REFERENCES

- Aist-Sagara, L., J. F. Cooper, R. E. McGuire, R. Parthasarathy, and M. Peredo, "Satellite Ephemeris Data Base for Magnetospheric Science Planning," *EOS Supp., Trans., AGU* 74, 270, 1993.
- Convery, P. D., J. Raeder, R. J., Walker, M. Ashour-Abdalla, "Magnetic Field and Flow Signatures at the Deep Tail Magnetopause," *EOS Supp., Trans., AGU* 73, 471, 1992.
- Fairfield, D. H., "Average and unusual locations of the earth's magnetopause and bow shock," *J. Geophys. Res.*, 76, 6700, 1971.
- Fairfield, D. H., "A Statistical Determination of the Shape and Position of the Geomagnetic Neutral Sheet," *J. Geophys. Res.*, 85, 775, 1980.
- Goodrich, C. C., and J. Lyon, "MHD Simulation of the Windsock Effect," *EOS Supp., Trans., AGU*, 73, 461, 1992.
- Goodrich, C. C., and D. A. McNabb, "The Advanced Visualization Laboratory of the University of Maryland, College Park," *IEEE Computer Graphics*, 26, 188, 1992.
- Goodrich, C. C., D. A. McNabb, R. Kulkarni, M. Peredo, and J. Lyon, "Interactive Software Tools for the ISTP Project: Science Planning, Data Analysis and Model Interpretation," *EOS Supp., Trans., AGU*, 74, 265, 1993.
- Hones, E. W., Jr., R. D. Zwickl, T. A. Fritz, and S. J. Bame, "Structural and Dynamical Aspects of the Distant Magnetotail Determined from the ISEE-3 Plasma Measurements, Planet. Space Sci, 34, 899, 1986.
- "IDL Reference Guide, Interactive Data Language Version 3.0," Research Systems, Inc., Boulder, CO, January 1993.
- Lyon, J., and J. A., Fedder, "Global MHD Simulations of the Aug 2-3, 1991 Campaign," *EOS Supp., Trans., AGU*, 73, 474 1992.
- McNabb, D. A., C. C. Goodrich, and M. Peredo, "Using AVS to further NASA Science Research and Mission Planning," in *Proc. 2nd Annual Intl. AVS User Group Conf.*, Lake Buena Vista, FL, IAC/NCSC, May 1993.

- Mish, W. H., M. Peredo, J. L. Green, and M. G. Reph, "ISTP Science Data Systems and Products," *Space Science Reviews*, 1994 (in press).
- Parthasarathy, R., M. Peredo, L. Sprayregen, and R. E. McGuire, "Multispacecraft Mission Planning/Analyses Resources at NSSDC/WDC-A-R&S," *Proceedings of the 1992 STEP Symposium*, 1994 (in press).
- Peredo, M., "Science Planning and Operations Facility," *STEP International Newsletter*, 3, August, 1993.
- Peredo, M., Ch-M. Wong, and H. K. Hills, "Data and Software Resources at the Satellite Situation Center," *Geofísica Internacional*, 31, 145, 1992a.
- Peredo, M., N. A. Tsyganenko, S. A. Curtis, D. P. Stern, and M. V. Malkov, "Solar Wind and IMF Control of the Magnetopause in a Tsyganenko Model with Tail Warping," *EOS Supp., Trans., AGU*, 73, 462, 1992b.
- Peredo, M., D. P. Stern, and N. A. Tsyganenko, "Are Existing Magnetospheric Models Excessively Stretched?," *J. Geophys. Res.*, 98, 15343, 1993.
- "PV-Wave Command Language Version 4.2 Technical Reference Manual," Visual Numerics, Inc., Boulder, CO, 1993.
- Raeder, J., "Ionospheric Convection Patterns for Northward IMF and Variable B_y : Simulation Results from a Coupled Global MHD Magnetosphere-Ionosphere Model," *EOS Supp., Trans., AGU*, 73, 465 (1992).
- SSC, "Satellite Situation Center Solar-Terrestrial Physics Planning Tool, Software Application Users' Guide, Version 2.0," NASA National Space Science Data Center, Greenbelt, MD, October 7, 1992.
- Sibeck, D., R. Lopez, and E. Roelof, "Solar Wind Control of the Magnetopause Shape, Location and Motion," *J. Geophys. Res.*, 96, 5489, 1991.
- Stern, D. P., and N. A. Tsyganenko, "Uses and Limitations of the Tsyganenko Magnetic Field Models," *EOS Trans., AGU*, 73, 46, 1992.
- Szuszczewicz, E.P., A. Mankofsky, P. Blanchard, C. Goodrich, D. McNabb, and D. Kamins, "SAVS: A Space and Atmospheric Visualization Science System," in *Visualization Techniques in Space and Atmospheric Sciences*, E. P. Szuszczewicz and J. Bredekamp editors, NASA Printing Office, (Washington, DC, 1994).
- Tsyganenko, N. A., "Global quantitative models of the geomagnetic field in the cislunar magnetosphere for different disturbance levels," *Planet. Space Sci.*, 35, 1347-1358, 1987.
- Tsyganenko, N. A., "A magnetospheric magnetic field model with a warped tail current sheet," *Planet. Space Sci.*, 37, 5-20, 1989.
- Tsyganenko, N. A., "Quantitative Models of the Magnetospheric Magnetic Field: Methods and Results," *Space Science Reviews*, 54, 75-186, 1990.
- Upton, C., T. Faulhaber, Jr., D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam, "The Application Visualization System: A Computational Environment for Scientific Visualization," *IEEE Computer Graphics and Applications*, pp. 30-42, July, 1989.
- Whipple, E. C., and H. Lancaster, Eds., "Toward a New Era of Global Solar-Terrestrial Research - The First IACG Campaign: Magnetotail Energy Flow and Non-Linear Dynamics," Science Applications International Corporation, October, 1992.

Visualization of the Electrostatic Potential Distribution in Both Polar Ionospheres Using Multiple Satellite Measurements

MARC R. HAIRSTON and RODERICK A. HEELIS

Center for Space Sciences, University of Texas at Dallas, Richardson, TX

FREDERICK J. RICH

Phillips Laboratory, Geophysics Directorate, Hanscom AFB, MA

During the time from December 1991 through March 1992, there were four operational DMSP satellites in polar orbit. All four satellites carried the SSIES plasma package which included an ion drift meter. Data from the drift meter, combined with the magnetic field data, allowed the calculation of the electrostatic potential in the ionosphere along the satellite's path. Simultaneous polar coverage by four satellites was unprecedented, providing researchers with almost continuous monitoring of the potential distribution in both hemispheres for the four month period. Combining the magnitude and location of the potential data from each of the four satellites in order to examine the varying potential distribution pattern in both hemispheres presented a major challenge in data visualization. The problem was solved by developing a three-dimensional presentation of the data where the potentials are color coded and represented by the vertical dimension. This paper presents examples from a computer animation of several days of data demonstrating evolution of the size and shape of the potential distribution, along with how these changes correspond to variations in other geophysical parameters, such as the IMF orientation and the K_p index.

1. INTRODUCTION

The National Center for Supercomputing Applications (NCSA) produced video "Numerical Modeling of a Severe Thunderstorm" [Wilhelmson et al., 1989] is perhaps the best known example of the application of computer animation to present a large and complex scientific data set in a clear manner. As such, it sets a standard that many scientists strive to equal in the animated visualizations of their own results. However, most scientists and researchers do not have access to the high-end computers and video resources necessary to produce such impressive animations. This paper will describe a system developed at the Center for Space Sciences at the University of Texas at Dallas for producing high quality stills and animations using relatively inexpensive equipment. More importantly, this paper will also demonstrate how these animations and visualizations are applied to the large data sets used in the analysis of the shape and evolution of the electrostatic potential distribution in the polar regions of the Earth's ionosphere.

2. BACKGROUND

The Center for Space Sciences builds an ion drift meter instrument as part of the Special Sensor-Ions, Electrons, Scintillation (SSIES) package, which flies on the polar-orbiting Air Force weather satellite series (the Defense Meteorological Satellite Program or DMSP) at an altitude of 800 km. This instrument measures the bulk ion velocities in the horizontal and vertical directions perpendicular to the satellite's velocity vector. The ion flows are sampled six times per second for each component and then averaged into four-second bins. As the satellite flies over the polar region, this four-second

resolution flow data is combined with the magnetic field data to calculate the electrostatic potential along the satellite's track. A more detailed description of the analysis techniques used here is given in Hairston and Heelis [1993] and Heelis and Hairston [1990]. The interaction of the interplanetary magnetic field in the solar wind as it moves past the Earth's magnetosphere generates an electric field across the surface of the magnetopause, which in turn produces a potential drop between the dawnside and the duskside. This potential drop is mapped down to the two polar ionospheres producing a potential difference there on the order of tens to hundreds of kilovolts. The exact distribution of the potential across the polar cap region is complicated and has been studied and modeled for the past twenty years [e.g. Heppner and Maynard, 1987; Heelis, 1984; and the references therein]. There is general agreement in these studies that when the interplanetary magnetic field (IMF) in the solar wind is pointed southward (i.e. $B_z < 0$) a region of positive potential forms on the dawnside of the polar region and a corresponding region of negative potential forms on the duskside. The shapes and sizes of these regions vary widely, but are believed to be influenced by the magnitude and sign of the y-component of the IMF. Here the y-axis is oriented parallel to the Earth's orbit with +y directed towards the duskside. When the IMF is oriented northward (i.e. $B_z > 0$), there is less agreement about the nature of the distribution [see Reiff and Burch, 1985] but it is generally accepted that the potential drop is smaller and the overall potential distribution pattern more complicated. While the potential perpendicular to the magnetic field in the polar region covers a two-dimensional area, the data from the satellite only covers a single slice of the pattern as the satellite crosses the pole. Thus, all of the modeling of the potential pattern to date has been based on averaging and binning of data from satellite passes at widely different times [e.g. Rich and Hairston, 1993; Heppner and Maynard, 1987].

The ideal situation would have multiple satellites in orbit simultaneously, each sampling a different region of the pattern. This would insure that a true representation of the global distribution of the potential could be constructed. Such an ideal situation occurred from December 1991 through the end of March 1992, when there were four opera-

tional DMSP satellites in orbit, each in a different local time orientation. This provided an unprecedented opportunity to map the overall electrostatic potential distribution simultaneously in both polar ionospheres and look at short term changes in these patterns. An initial analysis was conducted on a four-day time period from January 26 through January 30, 1992, a time period that also encompassed one of the extended Geospace Environment Modeling (GEM) campaign periods.

3. DESCRIPTION OF THE GRAPHICS

The overall ionospheric potential distribution in one hemisphere can be visualized as a deformed three-dimensional surface over the polar region (Figure 1). Here the magnitude of the potential at a given point is represented by the distance above or below a polar dial. The polar dial denotes the magnetic local times (MLT) and the circles of magnetic latitude down to 50 degrees. In this figure, the positive potential region is shown as the elevated ridge on the right while the negative potential region is shown as the bowl-shaped depression on the left. To further reinforce the distinction between the positive and negative regions, they are also color-coded using grey-white for positive and orange-brown for negative. The potential curve observed by a single satellite pass is presented in the figure as the bright white and orange line along the surface, thus giving a cross-section of the surface along that track. By visualizing the electrostatic potential distribution as this "drumhead" surface, the variations in the potential can be represented as variation in the size and shape of the "drumhead."

Since the potential data are available only along the satellite tracks, three-dimensional plots of the potential curves for all four satellites are overlaid for each hemisphere to form a rough, wireframe representation of the size and shape of the overall potential "drumhead." Figure 2 presents a typical frame from the animation showing an example of the data from both hemispheres during this period. Each frame represents a timestep of twenty minutes centered on the time given in the upper left corner. The northern and southern polar dials are presented as flat surfaces observed from the nightside looking towards the dayside. The polar dials show the most




Figure 1. This view shows the northern polar region of the Earth as seen from the nightside. The blue polar dial represents the magnetic latitude in increments of 10 degrees and the crosshair represents the magnetic local times of midnight, 6 am, noon, and 6 pm. The translucent surface represents the electrostatic potential distribution in the northern polar ionosphere with the magnitude of the potential at any point corresponding to the distance above or below the polar dial. Thus, the positive potential region forms a grey-white ridge on the dawnside and the negative potential region forms an orange-brown bowl-shaped depression on the duskside. A sample potential curve that would be observed by a single satellite pass is represented as the bright curve across the surface.

recent passes in each hemisphere for all four satellites. Since the orbital periods of all four satellites are around 105 minutes, this means the pattern in each hemisphere is composed of passes of differing ages. In order to differentiate between the most recent passes and the older passes, the potential curves are color-coded [Tuft, 1990]. A new pass is one in which the satellite reached its highest magnetic latitude during the twenty-minute period of the current frame, and is colored the brightest. A pass that occurred before the current frame, but within the past hour is given a medium color, and any pass more than one hour is shown in the darkest colors. Thus, in the animation the potential curve for a given pass in a given hemisphere will be bright when it first appears, fade to a

medium value for the next two frames, then fade to a dark tone for another two or three frames before it is replaced by the curve from that satellite's next polar pass in that hemisphere. The observed maximum and minimum potential on each satellite's pass are represented by vertical red lines between the curve and the polar dial. For times of southward IMF, these vertical lines roughly denote the polar cap boundary on each polar dial. In the lower left corner of the frame is a plot of the K_p index during the four-day period of this data set with a vertical bar marking the current time.

In the northern hemisphere in Figure 2, the four potential curves suggest a surface with a large dome-shaped positive potential region that extends from the dawnside past the noon-midnight line and

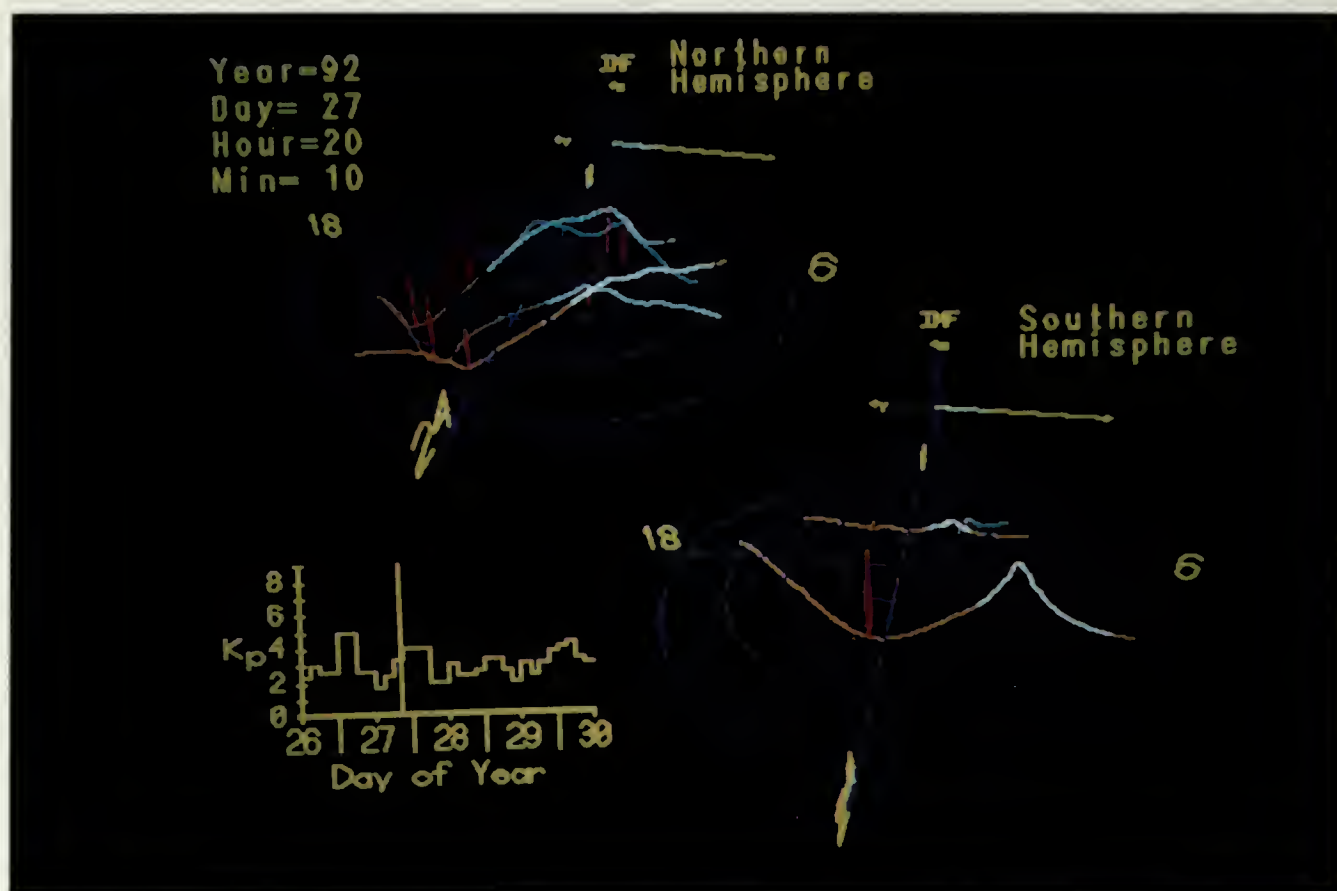


Figure 2. This is a typical frame from the animation showing the potential curves observed by the satellites in both hemispheres presented on polar dial graphs indicating the magnetic local times (MLT) and magnetic latitudes. Note how the curves mesh together to form a wireframe model of the surface describing the overall potential distribution. New passes that occurred during the 20 minute period represented by this frame are shown in brightest colors. Previous passes less than one hour old are shown in medium colors and passes more than an hour old are shown in dark colors. The vertical red bars indicate the location and magnitude of the maximum and minimum potentials observed along each pass. In the northern hemisphere the largest maximum bar represents 38.8 kV at 5.7 MLT and 76.9 degrees magnetic latitude, while the largest minimum bar represents -61.6 kV at 21 MLT and 69.1 degrees magnetic latitude. The yellow arrows on the crosshairs on the dayside of the polar dials represent the averaged magnitude and orientation of the IMF in the y-z plane during the time period for this frame. Here $B_y = -18.2$ nT and $B_z = -0.8$ nT. In the bottom left corner is a graph showing the K_p index for this study period with a vertical bar indicating the current time.

partway into the duskside. The negative potential region is shaped like a valley and confined to the duskside. The fact that the different colored curves merge to give a coherent shape indicates that this potential distribution had been steady for some time. This distribution is consistent with the observed distribution during times when the IMF B_z is weakly negative (southward) and B_y is strongly negative [Rich and Hairston, 1993; Heppner and Maynard, 1987]. At these times the ion convection pattern in the northern hemisphere is dominated by a large dawn cell. The averaged IMF orientation and magnitude in the y-z plane for this twenty minute period is given as an arrow in the crosshair on the dayside of the polar dials. Note that the

orientation of the crosshair is what an observer would see if looking towards the sun (i.e. +z is upward and +y is to the left). There is a time lag of about 36 minutes between the time the IMF was observed in the solar wind by the IMP-J satellite and the time it is displayed on the crosshair. This time lag corresponds to six minutes for the solar wind transit time between the IMP-J satellite and the nose of the Earth's magnetopause, plus a 30 minute delay for the entire polar ionosphere to respond to the IMF. This 30 minute response time of the ionosphere was determined empirically using this four-day data set.

In the southern hemisphere in Figure 2, the pattern formed by the three potential curves is quite

different. (The fourth curve in the 1400-1600 MLT region is flat showing that this satellite pass did not enter the auroral region.) Here the positive potential region is shaped like a ridge that forms a crescent shape on the dawnside while the negative potential region forms a large bowl-shaped basin extending from the duskside into the dawnside. This is also consistent with the observed potential distribution in the southern hemisphere during the times when the IMF is predominantly oriented in the -y direction and a large dusk cell dominates the ion convection pattern [Rich and Hairston, 1993; Heppner and Maynard, 1987]. Having multiple satellites providing simultaneous observations in both polar ionospheres clearly demonstrates that the potential distributions in both hemispheres are not mirror images of each other. It should be noted that the orientation of both the northern and southern hemispheric polar dials in Figure 2 match the

orientation seen in Figure 1 (dusk on the left, dawn on the right). This orientation was chosen to emphasize the comparisons between each side of the pattern and its counterpart in the other hemisphere.

4. ANALYSIS

Once all the frames were generated for this study period, they were connected together on the computer so they could be played in sequence as an animation. This proved to be an effective means of presenting the four days of data from four satellites in two hemispheres along with the IMF and K_p data in a form that was clear and understandable to the investigator. The interactive nature of the animations allowed the investigator to "play" the data at variable rates, to stop and examine a single frame at length, or even move backwards through the data. This animation presented the evolution of the

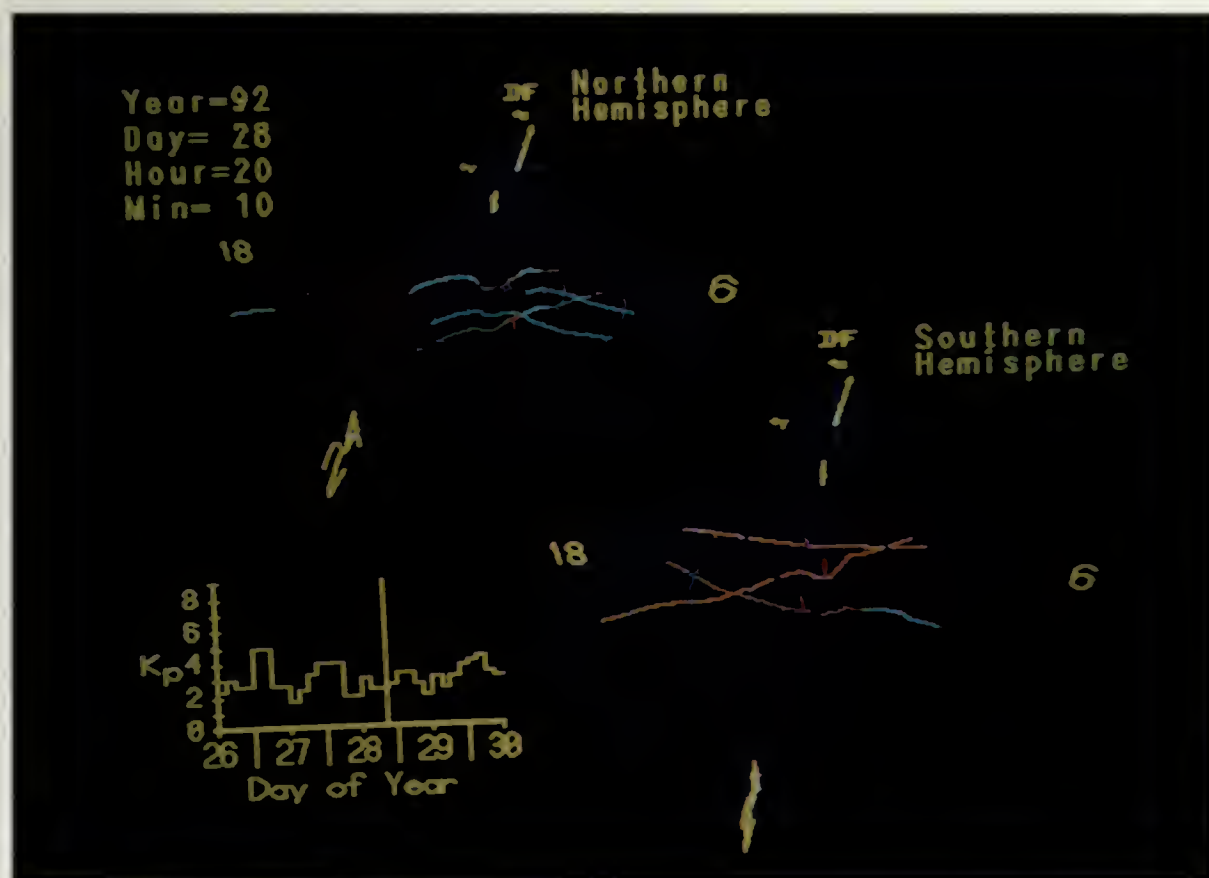


Figure 3. This frame presents a typical distribution seen when the B_z component of the IMF had been strongly positive (northward) for over an hour. During this twenty-minute period the averaged component values were $B_y = -1.5$ nT and $B_z = +4.9$ nT. Notice that the overall distribution is no longer organized into the two-cell pattern seen in Figure 2. The surface described by these curves is much flatter than in Figure 2 indicating that the electrostatic potential differences are much smaller.

potential distributions in both hemispheres over time and demonstrated periods of both extended stability and rapid change. This animation was recorded onto videotape for presentation at the spring 1993 AGU meeting (see note at the end of the article). This data set provided a wealth of information that is still being analyzed about the behavior of the ionosphere in response to the IMF.

In this short report, only four cases from the four day study period will be presented. Figure 3 shows a typical pattern observed when the IMF orientation is predominantly B_z positive (northward IMF). The overall potential drop between the dawnside and duskside is much smaller than when B_z is negative (southward IMF), thus causing the surface to appear much flatter than it did in Figure 2. Also, each of the potential curves shows

multiple oscillations rather than the simple sine wave curve seen in the curves in Figure 2. Taken together the curves indicate that the overall distribution is complex in shape and less organized than during times of southward IMF. This distribution is consistent with other observations during times of strongly northward IMF [Cummnock et al., 1992; Heppner and Maynard, 1987]. During conditions of northward IMF there is little or no reconnection between the Earth's magnetic field and the IMF, which accounts for the small potential drop between the dawn and duskside. The large scale two-cell convection pattern seen during southward IMF becomes distorted or breaks up into multiple smaller cells during extended periods of northward IMF. A satellite crossing through these multiple or distorted cells sees multiple reversals in the direc-

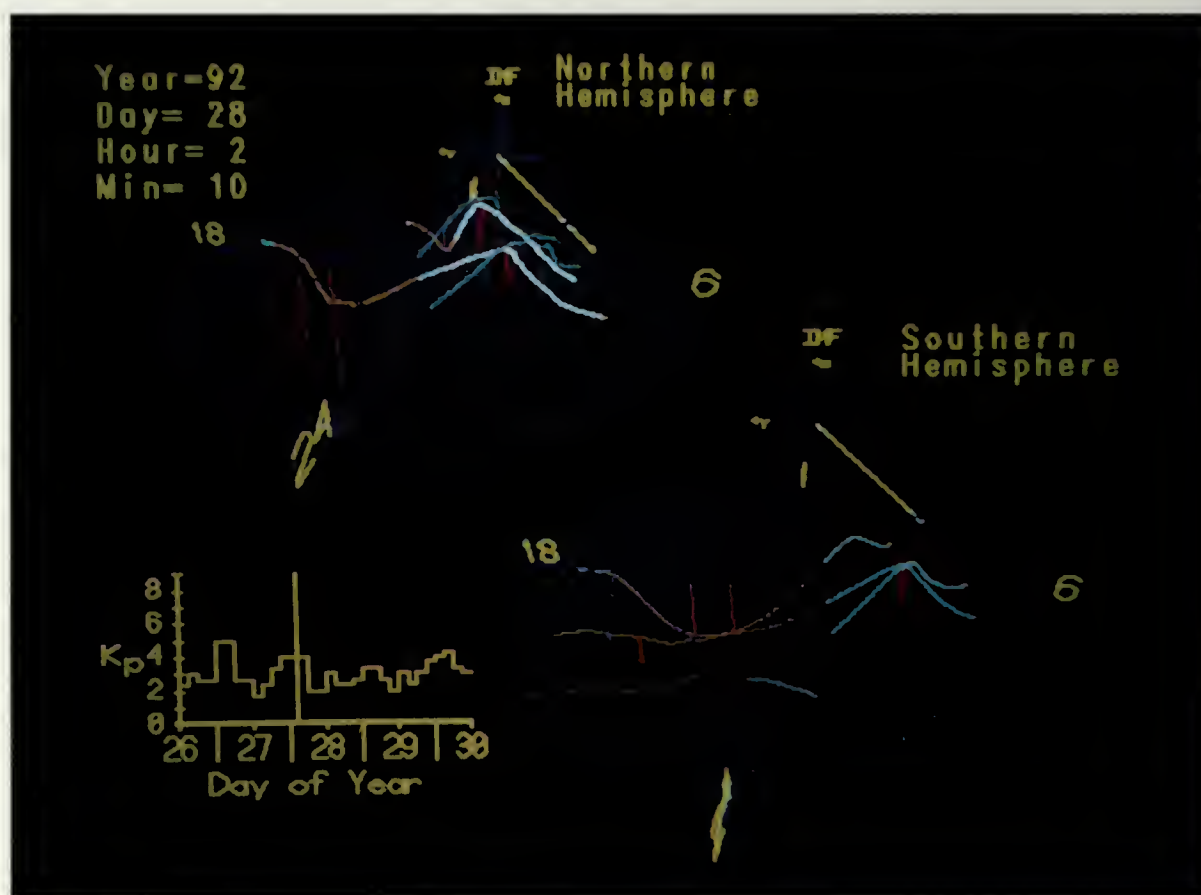


Figure 4. This frame presents a typical distribution when both B_y and B_z are negative and roughly equal in magnitude. During the twenty minutes of this frame the averaged IMF components were $B_y = -12.4$ nT and $B_z = -10.4$ nT. While the distribution is similar in overall form to the one in Figure 2, the rotation of the IMF towards the south changes the sizes of the potential regions. The northern positive potential region contracts back to the noon-midnight line compared to Figure 2, and the southern positive potential region increases in width. In the northern hemisphere the largest maximum bar represents 48.2 kV at 10.2 MLT and 75.9 degrees (magnetic latitude), while the largest minimum bar represents -62.1 kV at 19.3 MLT and 68.0 degrees.

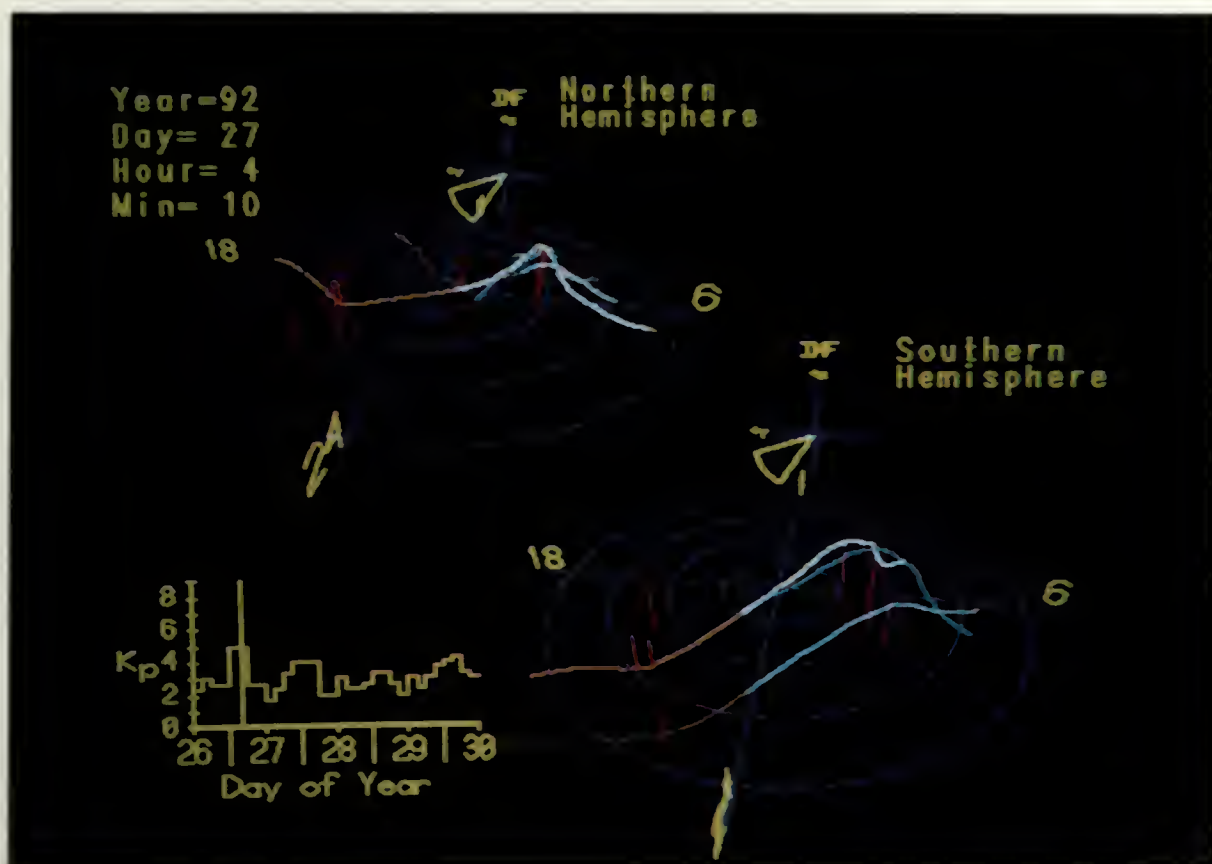


Figure 5. There were no available IMF data during the study period which showed the IMF orientation when B_y is positive and B_z is negative. However, this frame, taken from a period of the study when the IMF data were missing, presents a typical distribution pattern that is consistent other observations when B_y is positive and B_z is negative. The "wedge" on the IMF crosshair indicates an estimate of the orientation of the IMF. During such periods, the positive potential region in the northern hemisphere contracts over towards the dawnside and the positive potential region in the southern hemisphere expands out to past the noon-midnight line. In the southern hemisphere, the largest maximum bar represents 65.6 kV at 4.0 MLT and -74.2 degrees magnetic latitude, while the largest minimum bar represents -47.7 kV at 18.5 MLT and -67.6 degrees magnetic latitude.

tion of the horizontal ion flow during its pass, thus accounting for the multiple oscillations seen in the corresponding potential curve.

During times of predominantly negative B_z (southward IMF), the overall convection pattern becomes organized into two large cells: a dawn cell and a dusk cell. These convection cells correspond to the positive and negative potential regions (respectively) observed in the potential distribution. During times of southward IMF, there is a strong interaction between the Earth's field and the IMF, which drives the ion convection cells in the ionosphere and produces the potential drop between the dawn and duskside. Figure 4 shows the distribution for a frame when B_z and B_y had both been negative for several hours. Notice that the potential distribu-

tions are similar to those described in Figure 2. However, in this case, the magnitude of B_z is roughly the same as B_y which results in the widening of the crescent-shaped positive potential region in the southern hemisphere relative to Figure 2. Also, there is a corresponding shrinkage of the region of positive potential in the northern hemisphere to where it only reaches to the noon-midnight line [Rich and Hairston, 1993].

During times of southward IMF, it is the direction of the y-component of the IMF that largely determines the shape and extent of the convection cells and the overall shape of the potential distribution. Unfortunately, there were no IMF data available during this study period showing B_z negative and B_y positive. However, Figure 5 shows

a frame from a period without IMF data that corresponds to the distributions seen at other times when B_z was negative and B_y was positive. The positive potential region in the northern hemisphere has now contracted over to the dawnside forming a ridge shape while the negative potential region forms a shallow basin shape. In the southern hemisphere, the positive potential region has expanded past the noon-midnight line into duskside, confining the negative potential to a small region. Notice that the difference between the patterns in the two hemispheres is not as great as in the cases for Figures 2 and 4. This is consistent with the observations in Rich and Hairston [1993] for conditions when $|B_z| > |B_y|$. The wedge shape on the IMF crosshair indicates the absence of any actual IMF data, as well as the estimate that the

IMF orientation is in the negative B_z /positive B_y quadrant.

One effect of examining the data in an animated form is that the time history of the previous passes make changes in the potential distribution particularly easy to identify. Figure 6 presents a frame which shows a dramatic change in the potential distribution as a result of a change in the IMF. In the previous frames, B_z was positive and B_y was strongly negative. The faded passes in the northern hemisphere of this frame show that the positive potential region was weak and somewhat disorganized. The bright pass in the northern hemisphere is the only new pass to appear during this twenty minute period after B_z turned sharply negative. This new potential curve clearly shows that the positive region had increased in size and

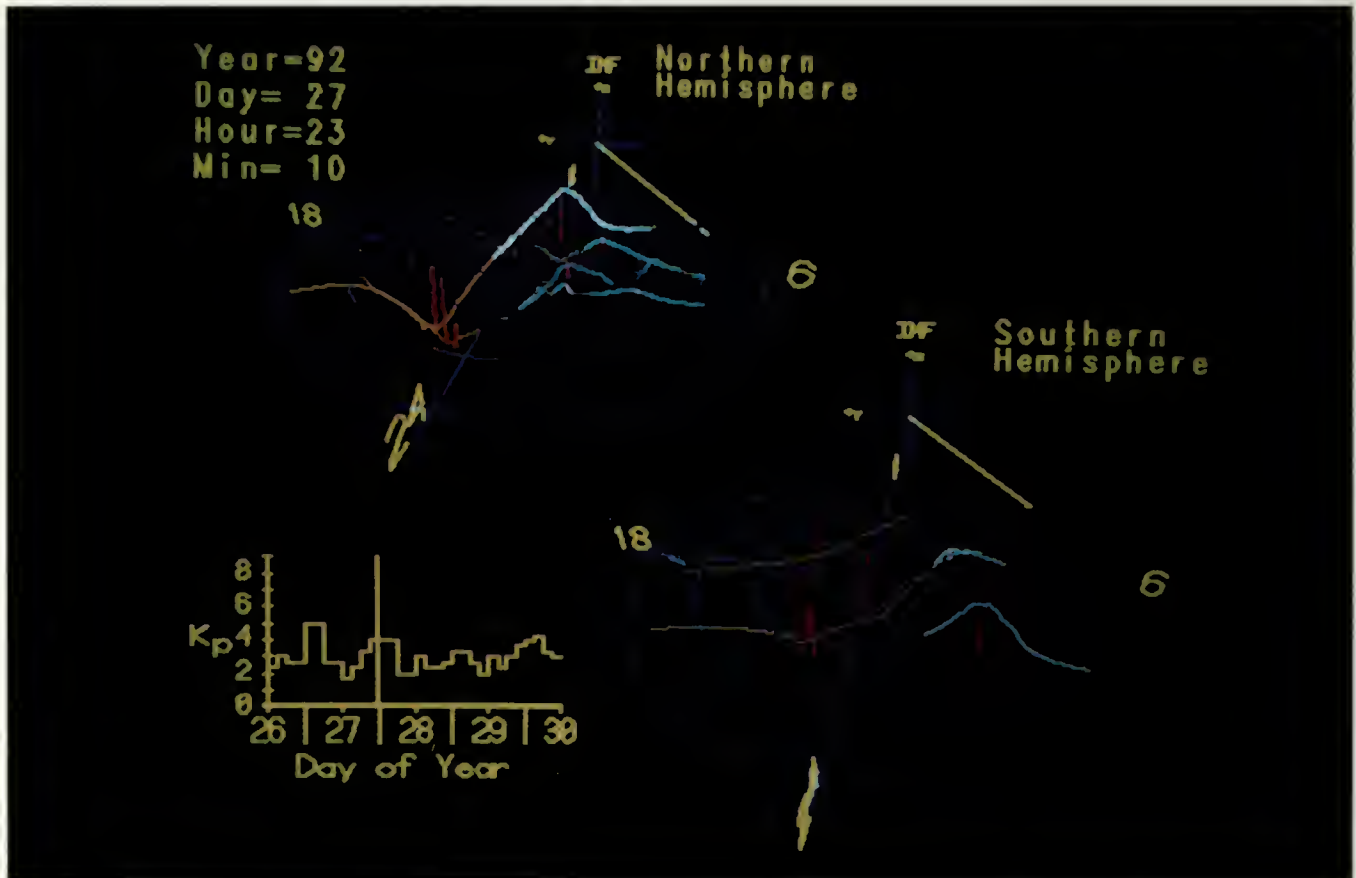


Figure 6. This frame demonstrates using the animation to identify changes in the potential distribution. The dimmed curves in the northern hemisphere are flat and disorganized in the positive potential region, while the negative potential region remains relatively deep and organized. This is indicative of the IMF orientation having a large negative B_y component and a large positive (northward) B_z component. During the previous twenty minute time period prior to this frame the averaged component values were $B_y = -8.7$ nT and $B_z = +6.3$ nT. However, the B_z component turned negative (southward) at the end of that period and during the current twenty minute period of this frame, the averaged component values have changed to $B_y = -14.4$ nT and $B_z = -9.3$ nT. The new potential curve that appears during this frame in the northern hemisphere shows the dramatic change in the size and shape of the potential distribution pattern.

strength compared to the previous potential curves in the northern hemisphere. Such events were used to quantify the response time of the ionosphere to various changes in the IMF as mentioned in section 3.

5. PRODUCTION OF THE ANIMATION

The animation and figures were produced using a VAX 4000 VLC workstation in conjunction with a Commodore Amiga 3000 personal computer. The processing of the raw telemetry data and the analysis to determine the electrostatic potential data for each pass were performed on the VAX. These processed data were used to generate individual frames in color on the VAX using a Tektronix plotting format. However, there was not an inexpensive way to make the VAX suitable for the playback of these individual frames as an animation. Even if it were possible, there was no inexpensive way to transfer such an animation from the VAX to video. This problem was solved by using the relatively inexpensive Amiga. The Amiga was hooked into the VAX cluster via ethernet and used a public domain terminal emulator called VLT [Weinstein, 1991]. The emulator enabled the Amiga to convert the plots downloaded from the VAX in Tektronix format to the IFF graphics format used by the Amiga. Once all the frames were downloaded and converted on the Amiga, commercially available software was used to string the frames together into an animation. This commercial software also allowed the user to modify the colors used and vary the playback speed of the animation. Since the Amiga was specifically designed as a video/multimedia platform, recording the video was a straightforward process of connecting a genlock (a device that converts the computer's RGB video output to standard NTSC video signal) to the computer, and then videotaping the output.

6. CONCLUSION

We have demonstrated in this paper that it is possible to create effective and useful scientific visualization animations using relatively inexpensive equipment. These animations enable researchers to see quickly and easily the shapes and sizes of the distributions of the electrostatic potential in the

ionosphere. While these images are not comparable to animations produced on high-end systems, the overall quality is still quite high and suits the needs of the researcher. Most of the work done to date in modeling the overall pattern of the electrostatic potential distribution has focused on static patterns which arise during long periods of stable IMF conditions. In reality, the IMF is rarely stable for time periods greater than an hour; thus the potential distribution in the ionosphere is quite variable and does not always match the static models. As researchers begin to model this variability and the time-dependence of the potential distributions, computer animations of the data, such as presented here will become an essential tool of that work.

Note: A copy of the animation of this work (NTSC only) will be sent to anyone who sends the authors a blank VHS videotape (still in the original wrapper) with a self-addressed stamped return envelope. Please mail requests to: Marc Hairston, Center for Space Sciences, University of Texas at Dallas, PO Box 830688 FO22, Richardson TX 75083-0688.

Acknowledgments. The authors wish to thank Dr. Ron Lepping of the NASA Goddard Space Flight Center and Drs. Barbara Emery and Gang Lu of the National Center for Atmospheric Research/High Altitude Observatory for providing the IMF data for this time period. The authors also wish to thank Professor Barry Treu of the Visual Arts Department of the University of Texas at Dallas for his assistance in producing the animation and Dr. Robin Coley of the Center for Space Sciences for his assistance in this work. This work was supported by the Air Force through Phillips Laboratory/Geophysics Directorate under contracts F19628-90-K-0002 and F19628-93-K-0007 and NASA grant NAGW-1665.

REFERENCES

- Cummnock, J.A., R.A. Heelis, and M.R. Hairston, Response of the ionospheric convection pattern to a rotation of the interplanetary magnetic field on January 14, 1988, *J. Geophys. Res.*, 97, 19449-19460, 1992.

- Hairston, M.R., and R.A. Heelis, High-latitude electric field studies using DMSP data, Tech. Rep., PL-TR-93-2036, Phillips Lab./Geophy. Directorate, Hanscom AFB, MA, 1993.
- Heelis, R.A., The effects of interplanetary magnetic field orientation on dayside high-latitude ionospheric convection, *J. Geophys. Res.*, 89, 2873-2880, 1984.
- Heelis, R.A., and M.R. Hairston, Studies of ionospheric dynamics utilizing data from DMSP, Tech. Rep., GL-TR-90-0047(I), Air Force Geophys. Lab., Hanscom AFB, MA, 1990.
- Heppner, J.P., and N.C. Maynard, Empirical high-latitude electric field models, *J. Geophys. Res.*, 92, 4467-4489, 1987.
- Reiff, P.H., and J.L. Burch, IMF B_y -dependent plasma flow and birkeland currents in the dayside magnetosphere, 2., A global model for northward and southward IMF, *J. Geophys. Res.*, 90, 1595-1609, 1985.
- Rich, F.J., and M. Hairston, Large-scale convection patterns observed by DMSP, *J. Geophys. Res.*, in press, 1993.
- Tufte, E.R., *Envisioning Information*, Graphic Press, Cheshire Connecticut, 1990.
- Weinstein, A., *A Valiant Little Terminal: A VLT User's Manual*, 121 pp., Stanford Linear Accelerator document SLAC-370 (Rev. #3), July 1991.
- Wilhelmson, R., H. Brooks, B. Jewett, C. Shaw, L. Wicker, "Study of a Numerically Modeled Severe Storm," video by the National Center for Supercomputing Applications, 1989.

Application of Advanced Computing Techniques to the Analysis and Display of Space Science Measurements

D.M. KLUMPAR, M.V. LAPOLLA*, B. HORBLIT*

Lockheed Palo Alto Research Laboratories,
Palo Alto, CA

A prototype system has been developed to aid the experimental space scientist in the display and analysis of spaceborne data acquired from direct measurement sensors in orbit. We explored the implementation of a rule-based environment for semi-automatic generation of visualizations that assist the domain scientist in exploring one's data. The goal has been to enable rapid generation of visualizations which enhance the scientist's ability to thoroughly mine his data. Transferring the task of visualization generation from the human programmer to the computer produced a rapid prototyping environment for visualizations. The visualization and analysis environment has been tested against a set of data obtained from the Hot Plasma Composition Experiment on the AMPTE/CCE satellite creating new visualizations which provided new insight into the data.

* Now at Meta Software, Inc., Campbell CA

1. INTRODUCTION

The use of increasingly sophisticated spaceborne instrumentation in space plasma physics generally has not been accompanied by application of increasingly sophisticated techniques for analyzing, assimilating, and displaying the observations. The task of creating effective visualizations of scientific data has, in the past, been time consuming and difficult. It has required significant programming and graphical design expertise and creativity, and results in visualizations with little flexibility. As data rates increase new techniques need to be put into practice to facilitate more rapid data analysis and interpretation. Advanced visualization techniques can manage the magnitude and complexity of the massive data sets generated by new instrumentation. We developed a prototype analysis and display environment to reduce the development time and increase the flexibility of the visualization process. A key requirement in the design of the system is that scientists have the capability to quickly change the way the data is viewed to permit different lines of inquiry.

One of the empirical space scientist's most challenging tasks is to unravel the effects of a large number of simultaneously acting physical processes using limited data gathered from space-based sensors. The scientist has a data set of limited scope from which to decipher cause-and-effect relationships and extract governing physical principles. Through visualizations the scientist creates pictures that facilitate his ability to explore the interrelationships in the data. By facilitating rapid generation of diverse visualizations, the scientists capability to uncover hidden relationships is greatly enhanced.

For most practicing instrumentalists, the process by which their one-of-a-kind pictorial representations are conceived and developed has not substantially changed in the last quarter century. The graphical design process consists of converting a scientific objective, observation, or concept into specialized software which produces the needed graphics. The process for creating the visualization software is labor intensive entailing man-weeks to man-months of collaboration between the programmer and the scientist. This enormous effort produces a single, inflexible visualization that often will be used to display the entire mission data set. Furthermore, these visualizations are produced one frame at a time generally without the benefits of animation. Typically, the space scientist looks at printed hard copies of each two-dimensional visualization to search for relationships. Such inflexible visualizations, combined with the labor intensive method used to create them, severely limit the scientist's ability to truly explore and analyze the data.

The Oakwood system presented here, addresses these issues by combining a cooperative computer-aided design (CCAD) [Friedell et al., 1991] approach with a domain specific visualization rule base to aid a scientist in generating visualization programs. The Oakwood system is conceptually based on the Lockheed Integrated Visualization Environment (LIVE), a cooperative, graphical rule-based environment for designing software, process, algorithm, and application visualizations [LaPolla 1992, LaPolla et al., 1993]. Oakwood is a visualization designer that utilizes commercially available analysis and visualization software packages as its foundation. The prototype system was implemented over Advanced Visual System's AVS visualization package. The Oakwood system allows a space scientist to easily express his analysis in terms of mathematics customary to the field. From this expression, Oakwood creates the appropriate visualization program, as well as the appropriate mathematical treatment. The visualization and treatment are then applied to space science data sets, facilitating analysis and producing three-dimensional color pictures.

In this paper, we describe the prototype development system and briefly illustrate its functionality. Its application to the space sciences domain has

been demonstrated by applying it to data obtained from the Hot Plasma Composition Experiment (HPCE) on the Charge Composition Explorer (CCE) satellite of the Active Magnetospheric Particle Tracer Experiment mission (AMPTE). In the following sections, we describe the specific space science domain problem (section 2), the prototype visualization system (section 3), and its application to the AMPTE data (section 4). We conclude with discussions of related work and describe the future direction of our work.

2. SPECIFIC SPACE SCIENCE DOMAIN SETTING

While Oakwood is adaptable to the analysis and display needs of many diverse disciplines, we have emphasized its initial application to the space sciences domain. More specifically, we have directed our attention to the arena where direct *in-situ* measurements of plasmas, fields, and corpuscular radiation are performed aboard spacecraft. This domain has special requirements and needs distinct from the broader space sciences domain where remote imaging is the primary tool. In comparison with the numerous efforts to advance the state of analysis and display for image processing, relatively little effort has been devoted to general processing and analysis tools for non-imaging space investigations characterized by direct measurement.

In imaging, one collects photons (or other quanta) arriving from a distant object to form a physical image. Spectral differentiation is often utilized to bring out specific physical characteristics of the scene and, thus, multi-spectral images are often collected. The primary characteristic of such images is that physical space is represented on the image plane. In contrast, in the direct measurements arena, the visualization and analysis parameter space is typically not described by physical dimensions, but rather by quantities which represent physical characteristics of the medium. Thus, direct measurement space-science data resides in an abstract space, rather than on a plane characterized by physical dimensions on an object. For example, a suite of direct measurement instruments might measure the D.C. magnetic field (a three-component field vector), the D.C. electric field (also a

three-component field vector), the flux of particles as a function of their energy of arrival, charge, and mass composition, and the A.C. electric and magnetic wave field spectra (energy density versus frequency). Each parameter is measured as a function of time and space (as the satellite moves along its trajectory) and all parameters vary in time and space with rates-of-change that are highly variable as the environment responds to externally applied forces, and as the spacecraft moves between different plasma regimes. The task of the analyst is to decipher the cause-and-effect relationships among the variations of the observed parameters and to determine the physical processes which govern the observed behavior. The directly measured parameters may be compared against one another using visualizations, such as line plots, scatter plots, and surface representations, or may be used to produce false images, where the orthogonal axes correspond to various parameters. The measurements may also be transformed into higher order quantities such as moments of the distribution function or phase space density. Other transformations utilize combinations of data from multiple sensors giving secondary parameters that represent global characteristics of the plasma (e.g. plasma beta, ratio of plasma frequency to gyrofrequency, etc.). The direct and the derived quantities are often analyzed in concert with predictions from theory. Thus the direct measurements area is a distinctly different enterprise from imaging and requires special analysis and visualization tools distinctly different from those employed in image processing.

3. THE PILOT SYSTEM

To use Oakwood, a space scientist first determines the type of analysis he wishes to perform. This analysis may be very specific, such as transformation of particle flux to phase space density. Or, depending upon the analysis in mind, he may use more general analysis techniques, such as Fourier transforms. The scientist enters his analysis into the system at the space science level via a menu of mathematical functions and formulae. Since the main goal is an empirical scientific analysis of large quantities of data, Oakwood's graphical user interface is oriented towards mathematics and data manipulation. The analysis choices are used to

create a calculation which is a combination of AVS modules, PV-Wave, and C++ functions. Those same choices are used as constraints by the rule-base visualization generation component of Oakwood. Oakwood also creates a custom graphical user interface for the interactive manipulation of the data and its visualization.

Each mathematical function in Oakwood is associated with a type of visualization rule. This association is used to constrain the visualization planner's choices. The Oakwood planner fires only those rules which satisfy these constraints. Each rule in turn creates further constraints that must be satisfied. The former constraints we will call analysis constraints. The latter constraints, those generated by the visualization rules, we will call derived constraints. Each visualization rule is responsible for functionality in the underlying commercial packages. For example, AVS does not support animation. However, animation may be achieved by displaying successive slices of a volume. Animation may also be achieved by reading in successive data files and recording the graphical output. Thus, the animation can be achieved in several different ways using the AVS package. The animation rules know this and choose a particular method to achieve animation. Derived constraints are syntactic constraints based on data type output by a given module in the commercial packages. In the above example, depending on the type of slicer chosen, the system may generate images or polygons, thus requiring a different choice of AVS modules later in the creation of the visualization.

Once the scientist has chosen an analysis, he may modify the visualization construction by directly manipulating the visualization rules through the graphical level as shown in the large window (lower right window) of Figure 1. The graphical rules are divided into three components, following the Cell Model [LaPolla, 1992; LaPolla et al., 1993] as displayed in the window. The first is the data description component at the left side of the window. The data component models the dimension, complexity and type of data input into the system. The second is the data manipulation component. This component is responsible for manipulating the data using mathematical techniques and formulae, such as Fourier transform or

matrix inversion. The third component, shown at the right center of the window, is the graphical or data presentation component. This component is responsible for how the data is to be presented to the user, whether as animation, as a solid, as a vector plot, etc. At the space science level the user manipulates concepts from the space science domain. At the graphical level, the scientist manipulates concepts more closely associated with data and graphics. The choices the user has at the graphical level are conditioned by the choices made at the space science level.

The final choice a scientist may make concerns how the rules are applied by the visualization planning engine (See Figure 1, experiment scope panel of the main window). A scientist may choose six different visualization construction methods: exhaustive, first, biggest, smallest, most connected, and least connected. Each of these methods must

satisfy all analysis constraints, as well as all derived constraints. The **exhaustive** method means that Oakwood tries to create all variations of a visualization allowed by the chosen constraints. After all the visualizations have been created, the scientist may then pick the ones he wishes to use. The **first** method means that Oakwood creates the first visualization it can. The **biggest** method creates a single visualization, which uses as many functional modules as possible. The **smallest** method creates visualizations with as few functional modules as possible. The **most connected** and **least connected** methods are AVS specific. Basically, they are similar to **biggest** and **smallest**. The analyst may also bound his computation. That is, he can determine how much of the resources to apply to the creation of visualization and data analyses (See Figure 1, experiment bounds panel at lower right).

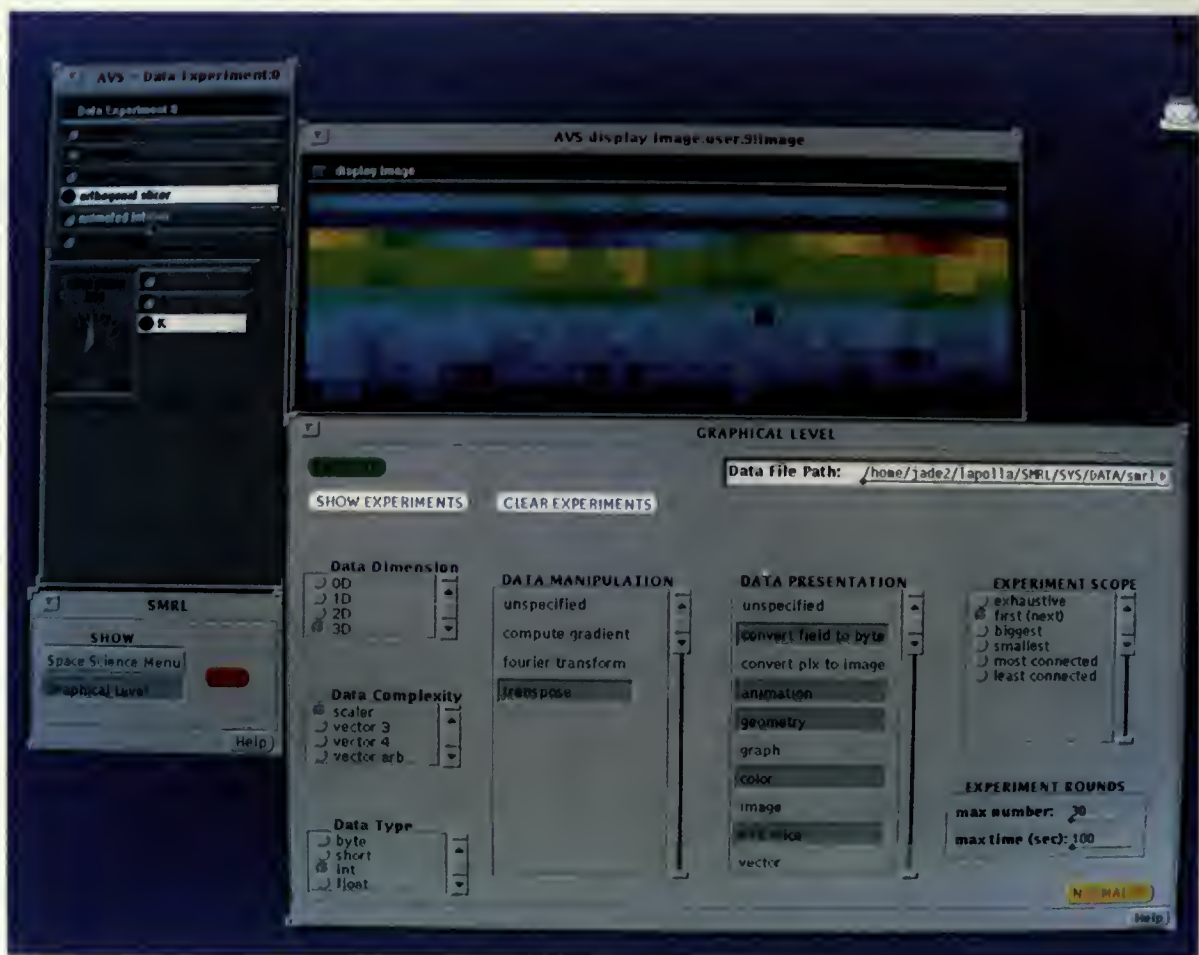


Figure 1. An example screen including windows for (1) the Graphical user interface (lower right); (2) selection of domain level or graphical level (lower left); (3) data experiment window (upper left); and (4) display image (upper right). The graphical level window is used to describe the data characteristics, manipulations, and to control the generation of images and their attributes as described in the text.

It is also possible to manipulate the visualization after Oakwood has created it. As Oakwood creates visualizations, it also creates graphical user interfaces for manipulating them. Since the Oakwood prototype uses AVS as its main driving package, it relies heavily on the AVS interface widgets such as "Vector Scale," "slice plane," "N Segment," or the choice of axis or interpolation methods. This is illustrated in Figure 2.

Once the visualization program and its interface have been created, the user selects data to be read in and Oakwood, through the commercial packages, produces the visualization (Figure 1). The upper

window of Figure 1 displays one spin of the AMPTE spectral data.

4. APPLICATION TO THE AMPTE DATA

The main strength of the Oakwood system is its ability to help the space scientist not only design and implement visualizations, but also design and implement programs to manipulate data. This section gives an example of the use of Oakwood and compares the current methods used by space scientists to the Oakwood methodology. The fact that the visualizations produced by the Oakwood system are sometimes simple and straightforward is unimportant. What is important is Oakwood's time-saving methodology and the freedom this imparts to the scientist. The rapid generation of diverse types of visualizations gives the scientist new freedom to roam through the data and to potentially uncover subtle but important interrelationships.

4.1 Data description

A subset of observations from the Hot Plasma Composition Experiment (HPCE) on the AMPTE/CCE spacecraft was chosen as a test application for the system. The Charge Composition Explorer (CCE) satellite of the AMPTE program was launched in August, 1984 into a near equatorial orbit (inclination 4.8° with apogee of $8.8 R_e$ and perigee 1108 km). The CCE was spin-stabilized at 10 rpm with its spin axis pointing approximately toward the sun. The spacecraft carried instrumentation to measure composition and charge state of ions over a very broad energy range, electrons, plasma waves, and the geomagnetic field. The data used here, from the HPCE, was taken by a set of eight magnetic electron spectrometers that measure the flux of electrons from 50 eV to 25 keV [Shelley et al., 1985]. Each spectrometer is operated at a fixed energy with an energy resolution of 50%. The eight instruments are co-aligned with their fields-of-view perpendicular to the spacecraft spin axis and are operated simultaneously, making measurements in unison every 155 msec. Thus, an eight point electron energy spectrum is obtained every 9.5° of satellite rotation. The spectrometers are collimated with a 5° full width conical field-of-

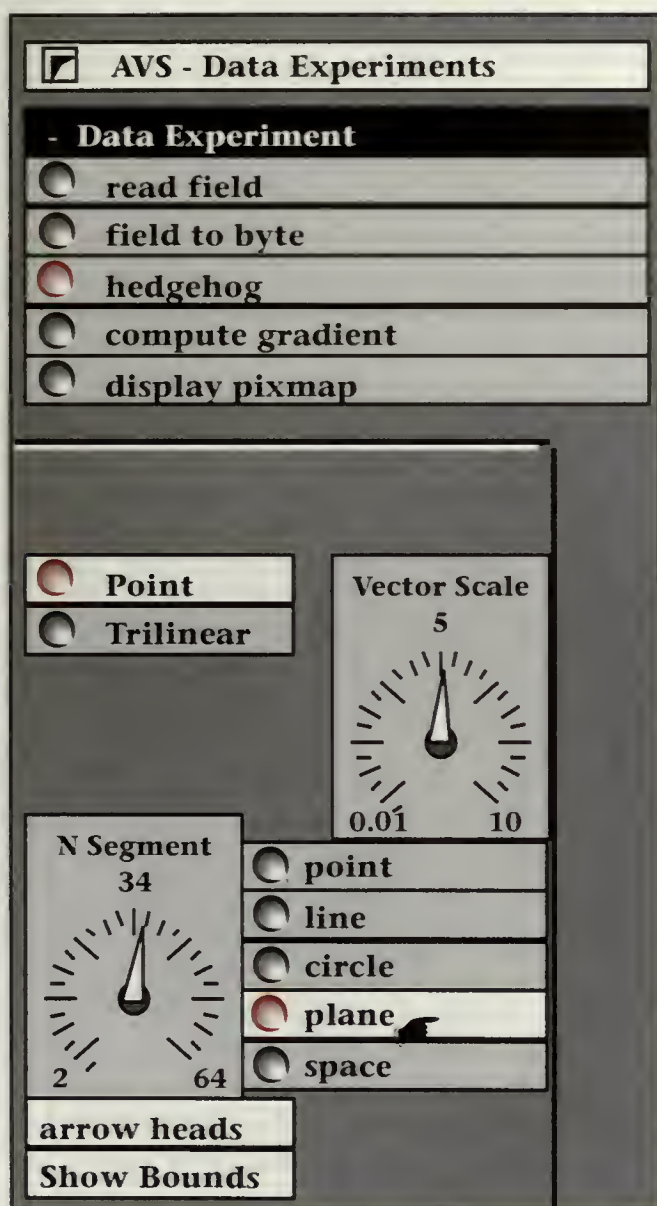


Figure 2. An example of the interface window for manipulating and modifying visualizations after creation of the prototype image.

view giving an effective full width field-of-view of $5^\circ \times 14.5^\circ$ during each measurement period. As the spacecraft rotates, the view directions sweep through a range of pitch angles the amplitude of which depends on the direction of **B** with respect to the spacecraft spin axis. The full pitch angle scan (0° - 180°) is achieved when **B** is perpendicular to the spin axis.

The output of the instrument may be viewed as a succession of time slices through energy space. Each time slice contains eight measures of the instantaneous electron flux (at the eight energies) and each successive slice is displaced by 9.5° of rotation in the spin plane of the satellite.

4.2 Current Visualization and Data Analysis Methods

It is customary practice to create specialized visualization templates, such as that shown in Figure 2 of Shelley et al. [1985], to display the essential data from an instrument. Several such unique visualizations are typically created to display the data for each instrument. The entire data run for the experiment is then typically dis-

played on these visualizations, creating thousands of frames of data which are stored on film, micro-film, or as hard copies. Selection of data for further analysis is carried out by visually scanning the hard copies or films to look for repeatable trends or interesting or unexplained signatures. Once event selection has taken place the analyst usually returns to the digital data for further analysis. This stage of investigation will involve further manipulation of the data and often results in the generation of additional specialized visualizations, which serve to display various characteristics of the events under analysis. Again the process is time consuming and labor intensive.

4.3 New Visualization

Figure 3 shows a visualization of the HPCE data described above. This visualization was created through PV-Wave and illustrates a simple, but effective enhancement to the presentation of particle spectrographic data. The basic display may be viewed as an energy vs. time spectrogram, a standard visualization in common use in space plasma physics. Since each time slice of eight-point




Figure 3. An example prototype visualization in energy-pitch angle-time space combines all three interrelated parameters of data sets commonly used in experimental space plasma physics (see text).

flux measurements represents a different pitch angle as the spacecraft rotates, co-display of the instantaneous pitch angle is necessary for full comprehension of the measurements. This essential parameter is incorporated into the plot of Figure 3 by allowing the instantaneous pitch angle to control the height of the three-dimensional surface. The corrugated sheet thus produced contains, on one panel, all of the relevant parameters in an easily assimilable representation. This new view replaces two separate pitch angle and spectrogram frames of the traditional display. With this visualization, the scientist can interactively choose to see one or several spins on the same visualization. Using the Oakwood system, the scientist can also see the same data as a traditional spectrogram and view multiple spins as animation. The power that Oakwood offers the scientist is the ability to quickly create and manipulate such flexible visualizations.

5. RELATED WORK

Oakwood is a hierarchical, cooperative design system for creating data experiments. Oakwood utilizes commercial packages and user provided code fragments as its mathematical and graphical backbone. This allows the user to quickly design and create mathematics programs and visualizations. Visualization production can be decomposed into articulation and design [Friedell et al., 1992, 1993; LaPolla et al., 1993]. Design generates an object conceptualization, which is an amalgam of primitive graphical sub-objects that are subject to geometric and topological relations (or constraints). The sub-objects, in this case, are functions in a commercial package organized using C++ classes and visualization rules. Articulation is the activity of providing a precise graphical description of an object conceptualization [Friedell et al., 1992, 1993]. The above is also true of data experiment production. Data experiment production can also be decomposed into articulation and design. For data experiment design in the Oakwood system, the object conceptualization and design discovery process is characterized as an amalgam of commercial package functions, a data characterization model, and design (search) strategies. Because of the multi-layered nature of Oakwood, there are not only geometric and topological constraints in the graphical model, and composition constraints in the mathematical model, but also

high-level contextual constraints imposed on the system by the application.

Various systems have been developed to support the articulation task for visualization in a general way, e.g., AVS, apE, and PV-Wave. The Visualization Design Environment, a component of LIVE, supports both articulation and design, but only for visualization generation.

Other systems support the articulation task for mathematics, e.g., MathMatica, PV-Wave. Few systems support the interaction between distinct graphical and mathematics packages, e.g., [Friedell et al., 1992, 1993]. Oakwood, however, supports both articulation and design for visualizations and mathematical programs. Because of Oakwood's open object-oriented architecture, design and articulation for other domains could also be supported.

We are aware of two related systems under development to provide analysis and visualization framework with application to the space sciences domains. LinkWinds, Linked Windows Interactive Data System, under development at Jet Propulsion Laboratory, has similar objectives to Oakwood [Berkin, private communication, 1993; Jacobson and Berkin, 1993] in that it is intended to provide a user-friendly environment to support rapid prototyping and execution of visualization applications. It differs substantially in approach, however, utilizing linkage of visuals through a graphical spreadsheet metaphor. Its scope extends beyond our plans for Oakwood by providing a multi-user environment for conducting cooperative research among remote sites. A second analysis and visualization tool known as SAVS is being developed jointly by Science Applications International Corporation (SAIC), Advanced Visual Systems, Inc. (AVS), and the University of Maryland. SAVS is a combined data acquisition, manipulation, analysis, and visualization system for application to solar-terrestrial and planetary sciences [Szuszczewicz et al., 1992]. Like Oakwood, this system is wrapped around the AVS visualization system, which provides a variety of tools for rendering volume data.

6. FUTURE DIRECTIONS

Several new NASA space flight missions, currently nearing completion of hardware development and scheduled to be launched in 1994, will carry new generation imaging particle spectrom-

eters that are expected to produce more comprehensive data on the plasma environment than any past missions. The Polar and Wind spacecraft of the GGS program and the Fast Auroral Snapshot Explorer of the Small Explorer program will each carry a comprehensive set of instruments to measure most of the important parameters governing the electrodynamic behavior of space plasmas and fields. The ultimate scientific success of these missions will depend upon the ability of the investigator teams to carry out comprehensive analysis of the data from the entire suite of instruments. It is missions like these that we have in mind as we develop the analysis and display tools described in this paper.

7. SUMMARY

We have described a prototype data analysis and visualization system that facilitates the rapid generation of visualizations with minimal user intervention. We have applied the system to the analysis and display of multi-parameter data collected by spaceborne instrumentation that measures the particle, plasma, and electromagnetic field environments through which satellites traverse. The system applies a rule-based approach implemented under the cooperative computer aided design paradigm to semi-automatically generate visualizations of a specified data set with minimal user intervention and guidance. When used under the control of a domain scientist, the system provides an environment for rapid creation of meaningful prototype visualizations. The flexibility to quickly explore potential interrelationships among one's data greatly enhances the scientist's ability to thoroughly explore his data.

Acknowledgments. This research has been supported by the Lockheed Independent Research Program. Support for the AMPTE program at Lockheed is provided under NASA contract NAS5-30565. The authors also acknowledge Evangelos Simoudis of the Lockheed Artificial Intelligence Center for many insightful discussions on the application of advanced computing techniques.

REFERENCES

- Friedell, M., S. Kochhar, M.V. LaPolla, and J. Marks, Integrated software, process, algorithm, and application visualization, *J. Visualization Computer Animation* 3, pp. 201-218, 1992.
- Friedell, M., M.V. LaPolla, S. Kochhar, S. Sistare, and J. Juda, Visualization of the behavior of massively parallel programs, *Supercomputing '91*, pp. 4572-80, Nov. 1991.
- Jacobson, A.S. and A.L. Berkin, LINKWINDS, A system for Interactive Scientific Data Analysis and Visualization, *High Performance Computing 1993*, Adrian Tentner ed., The Society for Computer Simulation, 1993.
- LaPolla, M.V., Towards a theory of abstraction and visualizations for debugging massively parallel programs, Hawaii Int'l. Conference on System Sciences-25, Jan. 1992.
- LaPolla, M.V., J. Sharnowski, B.C.H. Cheng, and K. Anderson, Data Parallel Program Visualizations form Formal Specifications, *Journal of Parallel and Distributed Computing* 18, 252-257, 1993.
- Shelley, E.G., A. Ghielmetti, E. Hertzberg, S.J. Battel, K. Altwegg-Von Burg, and H. Balsiger, The AMPTE CCE Hot Plasma Composition Experiment (HPCE), *IEEE Trans. on Geoscience and Remote Sensing*, GE-23, 241, 1985.
- Szuszczewicz, E.P., A. Mankofsky, P. Blanchard, C. Goodrich, D. McNabb, and D. Kamins, SAVS: A Space and Atmospheric Visualization Science System, in *Advanced Data Handling and Visualization Techniques in Space and Atmospheric Sciences*, E. P. Szuszczewicz and J. Bredekamp editors, NASA Printing Office, Washington, DC, 1994.

Data Management, Archiving, Visualization and Analysis of Space Physics Data

C. T. RUSSELL

*Institute of Geophysics and Planetary Physics and
Department of Earth and Space Sciences
University of California, Los Angeles, CA*

A series of programs for the visualization and analysis of space physics data has been developed at UCLA. In the course of those developments, a number of lessons have been learned regarding data management and data archiving, as well as data analysis. The issues now facing those wishing to develop such software, as well as the lessons learned, are reviewed. Modern media have eased many of the earlier problems of the physical volume required to store data, the speed of access and the permanence of the records. However, the ultimate longevity of these media is still a question of debate. Finally, while software development has become easier, cost is still a limiting factor in developing visualization and analysis software.

1. INTRODUCTION

A typical space physics investigation produces several gigabytes of information annually. The goal of the investigator who receives these data is to obtain the maximum amount of scientific understanding within a limited budget and a limited time frame. In general, the available financial resources are insufficient to complete the scientific analysis in the available time. Thus, to ensure that the data can be eventually analyzed, the original data must be archived, as well as the data products derived from the initial analysis. This archive should be permanent and secure since most data are unique and most likely will not be superseded in the foreseeable future. Yet at the same time, access must remain simple because the data might be needed at any time.

The space physics group at UCLA has faced these problems for nearly 30 years in handling data obtained on the ATS-1 and 6, OGO 1-6, Apollo 15 and 16, ISEE 1 and 2, Pioneer Venus Orbiter, and Galileo missions. In response to the continual problem of optimizing science return in the face of constrained budgets on these missions, we developed a data management system involving standard data formats, utility programs, and analysis programs that remained fixed from project to project to minimize the software development associated with each project. In 1979, we began the development of interactive graphics software as a further aid to increased scientific productivity. Also, during the mid-1980s the author served as chairman of the National Academy's Committee on Data Management and Computation (CODMAC) and NASA's Planetary Science Data Steering Group. We draw on the experience gained in this software effort and the deliberations of these two committees and their reports [Committee on Data Management and Computation, 1982; 1986; 1988] in writing this report.

It is not our purpose in this paper to describe the details of the software developed during this effort or to impress the reader with its capabilities. Rather, our purpose is to share our experience and lessons learned and assist future users and developers of similar efforts in this and allied disciplines.

2. DATA MANAGEMENT

When one begins to analyze large amounts of data from disparate sources, the need for standardization becomes quickly obvious. If one is to minimize reprogramming, standard data formats are a must. At UCLA we use a very simple structure, referred to as a flat file system, in which the data are stored in a binary file with time running down the first column and the various measurements made at that time in successive columns of the data matrix. This format is, by nature, best suited for data that can be stored as time series, but of essentially any dimensionality otherwise. The descriptions of the contents of the binary file are contained in a separate ASCII "header" file describing each column and its units, the original source of the data in the file, the creation date, and any explanatory text the creator desires to keep with the data set. New data arriving at UCLA are converted to this system before use. Exported data are converted to other formats as needed by the user if the user cannot use our file system.

Accompanying our standard file system are a set of utility programs that allow us to inspect and manipulate these files. In developing these utility programs, we have tried to make them both easy to use and flexible. By maximizing the use of a standard set of utility programs, we minimize the programming effort required for any new task, and we minimize software efforts by making maximum use of tested routines.

It is important that all data be catalogued. Moreover, the location of all data should be recorded, and metadata, the data that describe the contents of the files, should be solidly linked to the data products. Easy access to the catalog is as important as access to the data. Providing this access may require the development of additional software tools.

Over the last decade and a half, we have generally found that commercial programs were poorly suited to our needs. In particular, time series manipulation and display were generally poorly handled while contouring and the display of maps

or images were emphasized in these tools. Another problem with commercial software is that it is usually difficult to tailor such software to one's special requirements. Commercial software can also be expensive, but it is seldom as expensive as developing one's own special purpose software. Moreover, software development requires time, as well as money. Complex algorithms and the associated codes are not developed overnight. One solution to this problem is to collaborate with groups requiring or possessing similar codes, so that development costs are spread over several applications. Access to source code is usually essential, since no two scientific applications ever seem to have identical requirements.

A serious problem in developing a comprehensive data management system for a scientific (or commercial) application is that the computing systems improve at a rate comparable to that at which the software is developed. Hence, once the software is completed, the computer system is no longer state of the art and a new computer and consequent software adaptations and developments are required in order to stay competitive. Furthermore, operating systems and the windowing and display environment of workstations continue to evolve at a rapid rate. Over the last few years, our own group has evolved from graphic terminals to workstations running SUNVIEW, to OPENLOOK, and currently we are evaluating MOTIF. It is difficult but in this environment one must strive to be as device and operating system independent as possible.

Finally, one important lesson we have learned and one to which the CODMAC reports came back repeatedly is that software development must take place with the ultimate user closely involved. Software developed without such close coupling usually takes a wrong turn.

3. DATA ARCHIVING

The need for attention to the archiving of space science data has been receiving increasing attention in recent years as the realization struck that access to space was becoming less frequent, and not more frequent, with time. New missions in general complement or supplement old missions, not duplicate them. It is no longer possible to obtain approval to repeat an old mission simply because instrumentation has improved. Thus, most data obtained in space are unique. A corollary of this

observation is that if data are lost through some catastrophe, they are unlikely to be replaced in the foreseeable future.

Every group managing the sole or primary copy of a data set should establish an archive. Such an archive should have a searchable catalog, so that the contents of the archive can be determined readily. The data in the archive should be readily accessible. In the past, large data sets had to be stored in warehouses, but modern media have obviated the need for such warehouses. While the resources must be made available to copy the data, this copying cost is small compared to the cost of maintaining the warehouses. Care must be exercised in the choice of the new media, however. The lifetime of magnetic tapes of all types (of the order of 10 years) seems to be much less than that of optical disks and compact disks. Finally, we emphasize that with modern electronic networks, the site of an archive is becoming less important. It is possible for an archive to be distributed among several sites. For example, measurements of plasma temperature and density from a mission might be housed at one institution while the electric and magnetic field data are at another and the trajectory data somewhere else.

UCLA maintains an archive of magnetic field measurements from a number of past missions. In maintaining that archive we have learned a number of important lessons. First of all, we learned that seven and nine track tape does not provide a permanent archive. We still have much data on this medium, but we are attempting to copy those data as quickly as possible. At the urging of the National Space Science Data Center (NSSDC) we initially attempted to use 12-inch Write Once Read Many (WORM) optical platters. We copied all our experiment data records and supplementary data records for both the ISEE and Pioneer Venus projects from tape to WORM. However, we found the platters difficult to write and the equipment difficult to maintain. It is now clearly obsolete technology.

About two years ago, we purchased several Sony writers for erasable optical disks (EOD). Each have a data volume of about 600 MB. These were an immediate success among those analyzing and processing data. They are relatively inexpensive and are easy to use. Most recently we have begun to use the new 1.3 GB EODs. We also use EXABYTE 8mm tapes for backing up hard disks and for transporting data and have had no problems

in these applications. We are now beginning to write CD-ROMs. While the cost of CD-ROM media is less than that of an EOD (\$25 versus \$100 for a 600 MB disk), the CD-ROM writer and software is much more expensive (\$10,000 versus \$3,000). On the other hand, if one only wishes to read CD-ROMs and not to write them, the cost is trivial. Moreover, jukeboxes allowing access to large numbers of CD-ROMs are in the offing. Thus, the CD-ROM may be the medium of choice for future large data libraries. Finally, in choosing media one should remember that the access time of the various media types vary and that access time could become a limiting factor in an application. Table 1 lists average seek times and maximum transfer rate for a number of current media types. As attractive as CD-ROMs are, they have several disadvantages, such as limited capacity and moderately slow transfer rates.

Table 1. Data Access Rates

Medium	Average Access Time	Maximum Transfer Rate
Current magnetic disks	~8 msec	~10.0 MBytes/sec
Current EODs (1.3 GB)	~36 msec	~1.6 MBytes/sec
Optimum 12" WORM	~180 msec	~0.7 MBytes/sec
Double speed CD-ROM	~200 msec	~0.3 MBytes/sec
8mm "Exabyte" (5 GB)	~1000 sec*	~0.5 MBytes/sec
4mm DAT (2 GB)	~400 sec*	~0.35 MBytes/sec
6520 BPI 9-track tape (140 MB)	~350 sec	~0.2 MBytes/sec

*Time to search through 50% of a tape based on vendor's published "Search Rate" of ~2.5 MBytes/s.

The choice of media is not the only issue facing the archivist. Security, longevity, and documentation are also critical issues. All irreplaceable data should be copied and stored in two physically well-separated sites as a precaution against the loss of the archive due to some catastrophe, such as a fire or flood. Frequently handled data should be duplicated and one copy held in reserve. Documentation should not become separated from the data it documents. It is best that the documentation be put into machine readable form and stored on the same volume as the data. Furthermore, software used in creating data should be stored on the same volume as the data created. While software is not a perfect source of documentation, it is far better than having no documentation.

4. DATA ANALYSIS

The ultimate purpose of managing a data set and maintaining an archive is to facilitate the analysis of the data. Our developments in this area have also taught us some important lessons. First, the entry into the program should be easy. At most, the user should have to remember the name of the program, but even that can be avoided with careful system design. Similarly, access to the data should be made trivial. A program should lead the user to the requisite files. Operations on the data should take seconds and not minutes in any interactive application. Otherwise, the creative process is interrupted. The program should communicate with the user in English and the instructions to the program should be intuitive, so as not to require prior knowledge. The program must provide the functionality required by the user and it must be flexible. Finally, a screen is not the ultimate output device. A user needs to save the final results of the analytical session. Options for saving the results include paper, disk, tape, etc.

While it is easy to write down the desired attributes of analysis software, it is much more difficult to achieve it because specialized software development is expensive and time consuming. Moreover, no agency wishes to pay for this software development. Standard grant sizes are too small to afford a full time code developer and computer science grants generally go to advanced development rather than to applications. Again, the only practical solution in most situations is to attempt to collaborate with analysts with similar needs or to adapt existing codes.

Fortunately, many problems that we once faced in interactive graphics analysis have disappeared. The speed of machines is now no longer a problem in most applications. Memory sizes are usually quite sufficient. Data storage is now relatively inexpensive and data access times for most devices are not a limitation. Moreover, modern electronic networks have facilitated the exchange of both data and software. In fact, the various local, national and international networks have become essential to the conduct of space physics research at the present time.

5. INTERACTIVE GRAPHICS ANALYSIS IN THE UCLA SPACE PHYSICS GROUP

Our first developments in interactive graphics analysis began in late 1979 and 1980 when we obtained access to a Tektronix 4014 graphics terminal.

Our main computers in the 1970s and early 1980s were a series of HP minicomputers, an HP21MX, an HP F1000 and an HP A900 in succession. The 4014 was too expensive for us to buy multiple units, but it was sufficient to develop software and techniques for interactive graphics analysis and it soon became oversubscribed by students and staff. Fortunately, HP soon developed microcomputers (HP150) and terminals (HP2623) that had graphics and text screens (two windows) and that were inexpensive. We could then expand our service to multiple users although we had to be careful not to overload our minicomputer systems. In the mid-1980s, we were able to purchase a microVAX with a VMS operating system and we converted our programs to run under VMS. At this point we were able to export our software to a number of different collaborating groups. Thus, early versions of our software are in the hands of various research groups, some modified for user-specific applications.

In the late 1980s, we modified our software to run under UNIX on SUN workstations. We initially used the Template Graphics Package that was also available under VMS. We experimented with SUNVIEW, but as it became increasingly available and capable, we switched to using the OPENLOOK tool kit for X-Windows. We continue to upgrade the functionality of the software. The analysis software has been used in our analysis of data from the ISEE, Pioneer Venus, AMPTE, and Galileo missions. Some of the interactive graphics programs we have developed are described below. The series of programs all have names ending in ANAL, which stands for analysis. The first letter indicates the type of analysis.

5.1 Banal

This program was our first attempt at interactive graphics analysis and was designed to provide the tools needed to analyze microphysical processes (waves, discontinuities, boundaries) as observed in the vector magnetic field, B [Russell, 1983]. Thus there were tools to rotate, filter, Fourier analyze, and display magnetometer data. The program enables the calculation of averages, principal axes, wave properties, and hodograms, as well as time series.

5.2 Tanal

When a more global analysis was needed, such as viewing the pile-up of the interplanetary magnetic field in front of an obstacle, or the comparison of the field with a model over some volume of

space, it was helpful to include the trajectory (T) with the vector data and to display the field as vectors along the trajectory. This analysis was facilitated with the TANAL program, which also allowed plots versus altitude, as well as versus time, and in projection or planes.

5.3 Manal

The analysis of the state of the magnetosphere using ground-based magnetograms (M) has always been difficult because of the difficulty in accessing those records. During the International Magnetospheric Study (IMS), a number of digital stations were established. We acquired those data, sorted them by time and developed MANAL to provide access to the records and their metadata. MANAL also enabled their display and manipulation. As machines became more powerful, it became feasible to include dynamic spectral analysis in the list of interactive techniques. This capability is now included in MANAL. Moreover, we can perform cross correlation analysis, such as dynamic coherence and dynamic phase differences.

5.4 Uanal

The Pioneer Venus project assembled a key parameter data set patterned after the Unified Abstract Data Sets (UADS) of the AE and DE projects. We kept this key parameter data set on line and developed the UANAL program to access, manipulate, and display it. UANAL can display time series or altitude profiles, can display data from any instrument, by itself, or overlaid on other data. The user can select scales and intervals. Data can be combined to create derived parameters and can be compared with models. Measurements can be made from the screen with a cursor.

6. CONCLUSIONS

Modern computers and their software environments provide almost unlimited capabilities for data analysis and visualization. Nevertheless, data management issues are still a concern, in particular, standardization and user involvement. Data archiving has been facilitated by modern media, but documentation, media permanence, and

protection from catastrophes are still important issues. Finally, and very importantly, cost is still the limiting factor in software development.

7. APPENDIX: THE UCLA FLAT FILE SYSTEM

The flat file system is a two-dimensional data base system developed by the Space Physics group in the Institute of Geophysics and Planetary Physics at UCLA. The initial design of the system was developed by N.E. Cline in the early 1980s and was adopted for the analysis of the magnetometer data for both the International Sun Earth Explorer program and the Pioneer Venus Orbiter program. The initial code ran on the group's HP minicomputers. C.R. Clauer, then at Stanford University, ported the code to the VAX/VMS environment. In doing this, he made several changes in the format, most notably the origin of the time word, so that his files and those created at UCLA are not interchangeable. The Stanford system migrated to the Space Research Institute in Graz where extensive work took place on the TANAL analysis program and some on the BANAL program. The Hungarian PDS node also runs this version of the flat file system. Since the Graz group used IBM compatible PC computers extensively, it developed a PC-based flat file system and adapted some of the ANAL tools to the PC environment.

Eventually, the UCLA Space Physics group was able to purchase a VAX/VMS system. Because of their extensive investment in software and data sets in their original flat file format, the UCLA Space Physics group did not adopt the Stanford system, but ported their own standard format to the VAX/VMS. This version was exported to the JPL Space Physics group and enhanced by D. Winterhalter. In the late 1980s, UNIX became the operating system of choice. The Planetary Data System node at UCLA, the Planetary Plasma Interactions Node, decided to convert the flat file system to the UNIX environment. They did so, but in the process decided to increase the number of header files so that these flat files again were incompatible with the earlier set. The Space Physics Group, again, because of their large investment in program and data sets, decided they could not accommodate this format change and remained

with their original format. The net result is that there are several flat file formats in use, all with common ancestry, but slightly incompatible. Fortunately, it is simple to convert from one to another.

Flat files are useful for storing data which exist as multiple records or rows, each consisting of one or more data fields or columns. In space physics applications, they are most often used for time series of real numbers, but the data fields can contain most common data types such as: floating point, integer, character, and logical. A UCLA "flat file" actually consists of two files, an ASCII character header file and a binary data file. The ASCII header describes the associated data file with information, such as record length, number of rows, creation date, description of each data field or column, and user-written text. There are four parts to a header file. Lines 1 to 6 contain a physical description of the file (name, creation date, record length, number of rows and number of columns). The section from line 7 to the abstract line contains descriptions of each of the columns (column number, variable name, units, source, data type, location). The abstract consists of two parts, a fixed form, keyword section and a free-form user written section. The fixed form section contains information such as start and stop times, owner, flags, orbit number, averaging interval, etc. The binary data files have a fixed record length and can be accessed randomly. The user is not confined to sequential access. Since numerical data are stored in binary format, the record processing is much faster than in character data. The user can specify the number of data buffers to optimize disk access.

One of the important assets of the flat file system is the existence of an extensive library of utility programs which act on the flat files. One of the most useful programs is the file lister and plotting program, FL, which can be used to inspect any flat file. This program allows you to do fast and sequential searches for a given time, to compute averages, maxima, minima and standard deviations, to apply masks to the data, to convert to ASCII format, to plot selected columns versus time or each other and to perform various other functions.

The UCLA flat file system has been used extensively by the UCLA Space Physics group. However, the true mark of its utility lies in the scientific analysis that it has enabled. Almost all the analysis performed by this group and the resulting papers have relied on the existence of this software. Over 1000 papers at meetings and in journals and books have been supported by this software.

Acknowledgments. The author would like to thank the many individuals who have contributed to the UCLA space physics group's data management and interactive graphics analysis programs, especially to Q. Chen, N. Cline, H. Herbert, B. Littlefield, G. Maclean, J. Tung and K. Yamasaki.

This work was supported over the years principally by the National Aeronautics and Space Administration as part of the data analysis of the ISEE-1 and 2, and Pioneer Venus magnetic fields investigation. Some support has also been provided by the National Science Foundation under research grant USE 91-55988 and by NASA under research grant NGT-40005.

8. REFERENCES

- Committee on Data Management and Computation, *Data Management and Computation, Volume 1. Issues and Recommendations*, National Academy Press, Washington, DC, 1982.
- Committee on Data Management and Computation, *Issues and Recommendations Associated with Distributed Computation and Data Management Systems for the Space Sciences*, National Academy Press, Washington, DC, 1986.
- Committee on Data Management and Computation, *Selected Issues in Space Science Data Management and Computation*, National Academy Press, Washington, DC, 1988.
- Russell, C.T. Interactive analysis of magnetic field data, *Adv. Space Res.* 2(7), 173-176, 1983.

Interplanetary Space Science Data Base and Access/Display Tool on the NSSDC Heliospheric CD-ROM

N.E. PAPITASHVILI* and J.H. KING

*NSSDC, NASA/Goddard Space Flight Center,
Greenbelt, MD.*

Over the years, the National Space Science Data Center has accumulated a rich archive of heliospheric, magnetospheric, and ionospheric data, as well as data from most other NASA-involved science disciplines. To facilitate access to and use of these data, NSSDC has begun to put selected data onto CD-ROMs. This paper describes one such CD-ROM, and the access and display software developed at NSSDC to support its use. The data on the CD-ROM consist primarily of hourly solar wind magnetic field and plasma data from many near-Earth spacecraft (OMNI) and deep space spacecraft (Voyagers, Pioneers, Helios, Pioneer Venus Orbiter). In addition, 5-minute resolution IMP-8 and ISEE-3 magnetic field and plasma data are also included. Data are stored in both ASCII and CDF formats.

* On leave from WDC-B2/Geophysical Center, Moscow, Russia.

1. INTRODUCTION

The National Space Science Data Center (NSSDC) has collected data obtained by spacecraft in the Earth's magnetosphere and interplanetary medium for over three decades. NSSDC has accumulated the largest collection worldwide of interplanetary magnetic field (IMF), solar wind plasma (SWP), magnetospheric field, and plasma measurement data. These data sets are accessible on magnetic tapes and partially 'on-line' via computer networks. Magnetic tapes take time to retrieve data and do not provide any visualization software. The 'on-line' access is available for the scientific community in developed countries only. Thus, the existing data sets ought to be placed on modern media (e.g. CD-ROM) with 'user-friendly' software to be distributed worldwide.

In 1992 NSSDC began to create a series of Space Physics CD-ROMs, starting with heliospheric IMF and SWP data sets from various spacecraft travelling through the solar system: several near-Earth spacecraft (OMNI), Helios 1 & 2, Voyager 1 & 2, Pioneer 10 & 11, and Pioneer Venus Orbiter. The OMNI data base (one of the most widely used in the space physics community) contains hourly averaged IMF and SWP data near Earth for 1963-1991. In addition, 5-minute average data from near-Earth spacecraft IMP-8 and ISEE-3 are included. The 5-minute IMP magnetic field data are from the interplanetary, magnetosheath, and magnetotail phases of the IMP orbit, while the 5-minute IMP plasma data are primarily interplanetary with a small admixture of magnetosheath data.

The collected heliospheric IMF and SWP data sets have been provided to NSSDC by Principal Investigators (PIs) of the corresponding instruments

onboard given spacecraft. The details of provided data sets are presented in [Couzens and King, 1986] and may be retrieved from the 'on-line' Coordinated Heliospheric Observations (COHO) directory in the NSSDC ANONYMOUS account (see below for NSSDC ordering and access information). These data sets have been created by PIs in different formats and time conventions, and a variety of coordinate systems have been used for data presentation and location of spacecraft. Therefore, we have made uniform the data formats and coordinate systems. Also we have transformed data as needed in order to provide data from different spacecraft in the same coordinate systems. Note in particular that two ASCII versions of the OMNI data are included on the CD-ROM, one the familiar native format (no position data, field and flow vectors in GSE coordinates) and the other in the standard format (Earth position data, field and flow vectors in RTN coordinates).

The hourly resolution data files contain, for any given spacecraft, the heliocentric position of the spacecraft (or planet, for planetocentric orbiting spacecraft), the magnetic field intensity and Cartesian components, and the solar wind plasma flow speed, density, temperature, and, for some spacecraft, flow direction angles. The 5-minute resolution files have basically the same parameters, except that geocentric spacecraft position is given. Further details on the parameters are given below and are included in on-disk documentation files. The CD-ROM and associated software described in

this paper may be acquired by request of NSSDC at NASA Goddard Space Flight Center, Code 633.4, Greenbelt, MD, 20771, or 301-286-6695 by phone, or REQUEST@NSSDCA.GSFC.NASA.GOV on Internet.

2. DEFINITIONS OF COORDINATE SYSTEMS

The NSSDC Heliospheric CD-ROM contains the trajectory positions of all spacecraft providing the measured data (Table 1). Trajectory positions of all spacecraft are given in Heliographic Inertial (HGI) Coordinate System. The HGI coordinates of the planets Earth and Venus have also been computed to provide a uniform disposition of the near-Earth and near-Venus spacecraft together with the deep space observations. The Geocentric Solar Ecliptic (GSE) and Venus Solar Orbital (VSO) Coordinate System have been used to locate the positions of corresponding spacecraft relative to Earth and Venus. Some of the following coordinate systems are described by Roy [1988].

Solar Ecliptic Coordinate System (SE). The SE system is a heliocentric coordinate system with the Z axis normal to the ecliptic plane. The X axis goes toward the first point of Aries (Vernal Equinox, i.e. to the Sun from the Earth at the first day of Spring), and the Y axis completes the right handed set.

Heliographic Inertial Coordinate System (HGI). The HGI system is a heliocentric coordinate system with the Z axis along the Sun's spin axis

Table 1. Spacecraft and Data Sets Coverage, Resolution, and Coordinate Systems

Spacecraft	Data Coverage	Resolution	IMF	SW Plasma	Location
OMNI	11/1963- 07/1991	hour	RTN & GSE	RTN & GSE*	HGI
IMP-8	10/1973 - 06/1991	5-min	GSE	GSE*	HGI & GSE
ISEE-3	06/1979 - 12/1983	5-min	GSE	none	HGI & GSE
Helios-1	12/1974 - 12/1980	hour	RTN & SSE	SSE*	HGI
Helios-2	01/1976 - 03/1980	hour	RTN & SSE	SSE*	HGI
Pioneer Venus	12/1978 - 08/1988	hour	RTN & VSO	VSO*	HGI & VSO
Pioneer-10	03/1972 - 12/1988	hour	RTN	no angles	HGI
Pioneer-11	04/1973 - 08/1992	hour	RTN	no angles	HGI
Voyager-1	09/1977 - 08/1981	hour	RTN	RTN	HGI
Voyager-2	08/1977 - 12/1990	hour	RTN	no angles	HGI

* The SW plasma flow longitude angle is measured from -X (GSE, SSE or VSO) to -Y, i.e. positive for rotation of the flow in the same sense as the rotation of the Sun.

(positive to the North) and the X and Y axes in the solar equatorial plane. The X axis goes along the intersection line between the ecliptic and solar equatorial planes outward from the Sun in the general direction of winter solstice. The Y axis completes the right handed set. The ecliptic (SE) longitude of the X direction, i.e. the ascending node of the solar equator, was 74.37 degrees in 1900 and increases by 1.4 degree/century. To further orient the HGI system, note that at the start of 1994 the Pioneer-10 spacecraft moving down the tail of heliosphere was at zero degree HGI longitude.

Geocentric Solar Ecliptic (GSE) Coordinate System. The GSE system is a geocentric coordinate system with the Z axis northward from the ecliptic plane and the X and Y axes in the ecliptic plane. The X axis points toward the Sun from the Earth, the Y axis completes the right handed set. Often the GSE is used to refer to: (a) spacecraft centered system, which would be more appropriately called Spacecraft Solar Ecliptic (SSE), for a spacecraft that is near the ecliptic plane; and (b) other planet-centered system, e.g. Venus Solar Orbital (VSO) Coordinate System, where the X and Y axes are in the Venus orbit plane that is inclined about 3 degrees from the ecliptic plane.

RTN Coordinate System. The RTN system is centered at a planet or spacecraft. The R axis is radially away from the Sun. The T axis is the cross product of the Sun spin vector (North directed) and the R axis. The N axis completes the right handed set. Note that when the Earth is near its extreme heliolatitude excursion (e.g. +7.25 deg in early September) and its motion is parallel to the Sun's equatorial plane, the GSE and geocentric RTN systems line up, except that $X = -R$, $Y = -T$, and $Z = N$.

3. ACCESS AND DISPLAY SOFTWARE

Each data set has been created as a flat ASCII file with the logical record similar to the well-known OMNI data base. All parameters have been quality controlled (format, units, range, average values when possible, missing data, etc.), corrected (if necessary) and written in a similar format when possible (e.g. the SWP velocity values have format F6.1 across all data sets). Missing values have been replaced by combination of digit '9' in accordance

with the format of given parameter (e.g. F6.1 is filled by 9999.9).

In addition to the ASCII versions of the files, CDF (Common Data Format) versions are also provided on the disk. These CDF files are for the convenience of users having CDF software. For further insight into CDF, see *NSSDC CDF User's Guide* [1994] or call NSSDC's CDF User Support Office (301-286-9884). The software described in the following paragraphs is MS-DOS based, and addresses only the ASCII versions of the data files.

'Menu driven' and 'user-friendly' access and display software has been developed to review data coverage, to display data numerically and graphically, and to create files (subsets of CD-ROM files) for migration to magnetic disk for further analysis. The Master Screen (Figure 1) provides a single selection of spacecraft, shows general documentation, and allows a user to enter into the Intercomparison Software Package. This Package provides the ability to compare 'stack-plot' graphics of data from different spacecraft. The software will run in the 'hourly' or '5-minute' modes according to the data set selected.

The 'hourly' software works with any hourly resolution data set and offers the following options (see example on Figure 2):

- (a) display the Data Format Documentation File and provide a copy of this documentation file to the user's output file;
- (b) display the Monthly Data Availability Information Table, i.e. the 'up-to-one-month' screen (Figure 3) shows how much data are available for a given spacecraft for that month (100% of available data is denoted as Y, 0% - N, the numbers from 0 to 9 stand incrementally for each 10% of the data availability for a given day);
- (c) show simultaneously on the same screen the Daily Data Availability Information Line (bottom panel on Figure 3) which describes the data availability (Y or N) for each UT hour of a given day;
- (d) display the data listing on the screen (Figure 4);
- (e) create an ASCII file of the selected parameters and given time interval and display these parameters graphically;

- (f) display data for an arbitrary number of days, from 1 day (24 hours) to 31 days (Figure 5); and
- (g) display a scatter plot, i.e. to show the dependence of one parameter on another (Figure 6).

Note that the default plotting action is to plot time series about the mean value of the parameter for the interval, and to show the scale via the standard deviation measure to the left of the ordinate. Through the use of 'L-Change level' and 'S-Change scale' options in the display software, the user may fix the center line and the scale to meet his needs or taste.

The '5-minute' software provides similar options as the 'hourly' software. Some different display capabilities of this 5-minute mode are shown in Figure 7 (daily display) and Figure 8 (crossing day boundaries). The Intercomparison Software Package shows the 'spacecraft menu' which provides the ability to select files of corresponding data (Figure 9) from multiple sources. The program merges these files and shows a 'stack-plot' graphical display of selected parameters for comparison. There is an option that allows comparison of the 'hourly' and '5-minute' data on the same screen (Figure 10).

The hardware/software environment needed for this application software is MS-DOS, 1 MB of usable memory, and any IBM-PC compatible monitor, although color monitors will have their color capability used by the software. Other systems (UNIX, etc.) can access the ASCII files on the CD-ROM from their environments, although we provide no specific software support for such

systems. Other systems with CDF software can access the CDF files on the CD-ROM; see earlier comments about CDF.

4. CONCLUSION

The 'user-friendly' approach and easy access to the interplanetary data, developed for the first NSSDC heliospheric CD-ROM, is consistent with the concept of providing a Personal Data Center on CD-ROM. The wide distribution and use of IBM-PC compatible computers and CD-ROM readers provide users easy access to the data previously stored in the World Data Center archives. This effort brings the World Data Center data bases to scientists in any corner of the world.

Acknowledgments. We thank Principal Investigators and their colleagues who initially provided the data to NSSDC. We also thank J. Cooper and R. Parthasarathy for their contribution to the preparation of CD-ROM.

REFERENCES

- Couzens, D.A., and J.H. King, *Interplanetary Medium Data Book - Supplement 3, NSSDC/WDC-A-R&S, 86-04*, 1986.
- Roy, A.E., *Orbital Motion*, Adam Hilger Publ., Philadelphia, U.S.A., 552 pp., 1988.
- NSSDC CDF User's Guide, Version 2.4*, NASA/Goddard Space Flight Center, Greenbelt, MD, February 1994.



Figure 1. Master Screen of the NSSDC Heliospheric CD-ROM.



Figure 2. An example of the 'hourly based' software menu.

VOYAGER-1 DATA AVAILABILITY INFORMATION TABLE							
DAYS:	05	10	15	20	25	30	AVER.
IMF B Scalar, nT	Y Y 8 8 7 8 7	Y 6 8 4 6 6 7 8 8 8 8 8 7 7 9 Y 8 8 7 7 9 7 8 7					81%
IMF BR, nT (RTN)	Y Y 8 8 7 8 7	Y 6 8 4 6 6 7 8 8 8 8 8 7 7 9 Y 8 8 7 7 9 7 8 7					81%
SW Plasma Speed, km/s	Y Y 7 8 7 8 7	Y 6 8 4 6 6 7 8 8 8 8 8 7 7 9 Y 8 8 7 7 9 7 8 7					81%
SW Plasma Density, N/cm ³	Y Y 7 8 7 8 7	Y 6 8 4 6 6 7 8 8 8 8 8 7 7 9 Y 8 8 7 7 9 7 8 7					81%
SW Plasma Temperature, K	Y Y 7 8 7 8 7	Y 6 8 4 6 6 7 8 8 8 8 8 7 7 9 Y 8 8 7 7 9 7 8 7					81%
0 for 0-10%, 1 for 10-20%, ..., 9 for 90-100%, N - NO DATA, Y - ALL DATA							
DAY=10	SW Plasma Density, N/cm ³						
HOURS:	04	08	12	16	20	24	
	Y Y Y Y Y Y Y Y	Y Y Y Y Y Y Y Y	Y Y Y Y Y Y Y Y	Y Y Y Y Y Y Y Y	N N N N N Y Y Y		
Use Up/Down Right/Left Arrows ENTER - Monthly Listing ESC - Exit							

Figure 3. Data Availability Information Table for 'hourly' data sets.

	YEAR=80	MONTH=1	DAY=10	VOYAGER-1 HOURLY VALUES:				
IMF B Scalar, nT								
0hr	0.70	0.68	0.59	0.60	0.63	0.55	0.57	0.61
8hr	0.62	0.60	0.59	0.56	0.54	0.53	0.57	0.68
16hr	0.73	999.99	999.99	999.99	999.99	0.51	0.49	0.51
IMF BR, nT (RTN)								
0hr	0.11	0.02	-0.12	-0.29	-0.40	-0.19	-0.06	0.02
8hr	0.03	0.11	0.14	-0.02	-0.28	-0.29	-0.00	0.28
16hr	0.33	999.99	999.99	999.99	999.99	0.09	0.14	0.27
SW Plasma Speed, km/s								
0hr	380.1	376.3	375.0	375.0	376.4	373.2	376.7	378.8
8hr	377.4	375.8	375.7	378.3	370.5	369.2	371.7	371.8
16hr	370.8	9999.9	9999.9	9999.9	9999.9	368.9	368.8	367.5
SW Plasma Density, N/cm ³								
0hr	0.2153	0.2099	0.2933	0.2633	0.2827	0.3425	0.3149	0.2914
8hr	0.2810	0.3147	0.3337	0.2723	0.2872	0.3316	0.2209	0.1532
16hr	0.1166	99.9999	99.9999	99.9999	99.9999	0.2214	0.2221	0.2257
SW Plasma Temperature, K								
0hr	9005.	8424.	8567.	7589.	8141.	9912.	7589.	6670.
8hr	6418.	6172.	6798.	9453.	14348.	11026.	15878.	19602.
16hr	17076.	9999999.	9999999.	9999999.	9999999.	24443.	26175.	29549.
ENTER-Continue PgUp-Back Screen G-Graphics HOME-Main Menu ESC-Exit								

Figure 4. Listing of hourly mean values of selected parameters.




Figure 5. The 'multi-day' presentation of 'hourly' mean data.




Figure 6. An example of the 'scatter plot' option.




Figure 7. An example of 'daily' variation of '5-minute' data.




Figure 8. An example of the 'crossing day boundaries' plots.




Figure 9. Intercomparison of data from different spacecraft.




Figure 10. Intercomparison of 'hourly' and '5-minute' data.

SAVS: A Space and Atmospheric Visualization Science System

**E.P. SZUSZCZEWICZ, A. MANKOFSKY,
P. BLANCHARD**

Laboratory for Atmospheric and Space Sciences, Science Applications International Corporation, McLean, VA

C. GOODRICH and D. McNABB

University of Maryland, College Park, MD

D. KAMINS

Advanced Visual Systems, Inc., Waltham, MA

The research environment faced by space and atmospheric scientists in the 1990s is characterized by unprecedented volumes of new data, by ever-increasing repositories of unexploited mission files, and by the widespread use of empirical and large-scale computational models needed for the synthesis of understanding across data sets and discipline boundaries. The effective analysis and interpretation of such massive amounts of information have become the subjects of legitimate concern. With SAVS (a Space and Atmospheric Visualization Science System), we address these issues by creating a "push-button" software environment that mimics the logical scientific processes in data acquisition, reduction, and analysis without requiring a detailed understanding of the methods, networks, and modules that link the tools and effectively execute the functions. SAVS provides: (1) a customizable framework for accessing a powerful set of visualization tools based on the popular AVS visualization software with hooks to PV-Wave and access to Khoros modules, (2) a set of mathematical and statistical tools, (3) an extensible library of discipline-specific functions and models (e.g., MSIS, IRI, Feldstein Oval, IGRF, satellite tracking with CADRE-3, etc.), and (4) capabilities for local and remote data base access. The system treats scalar, vector, and image data, and runs on most common Unix workstations. We present a description of SAVS and its components, followed by several applications based on generic research interests in interplanetary and magnetospheric physics (IMP/ISTP), active experiments in space (CRRES), and mission planning focused on the Earth's thermospheric, ionospheric, and mesospheric domains (TIMED).

1. INTRODUCTION

1.1 An Integrated System

SAVS represents an integrated system concept with a "push-button" approach to the access and implementation of tools necessary to:

- search for, acquire, and read data;
- establish satellite tracks, their relevant parameters, and their co-registrations with allied multi-spacecraft and ground-based diagnostics;
- interactively interpret and analyze the findings;
- compare the results with mature and/or developing models; and
- creatively visualize the overall process and resulting scientific products.

Much of the SAVS effort has been devoted to the development and implementation of a user-friendly architecture focused on ease and functionality. Within our "push-button" approach, we have embedded a logical interactive data handling and analysis methodology, a mathematical toolbox, an interactive interpreter, and the AVS visualization software (Upson et al., 1989). As part of our effort, we are customizing AVS, its tools, and its operation for the NASA scientific user environment. Within this framework, we are creating custom functions, and integrating local and remote data base access, data translation, and formatting tools in order to provide scientific data manipulation, analysis, and display capabilities not available in the basic AVS system.

1.2 Uniqueness of the Architecture and Its End-to-End Approach

The SAVS architecture recognizes that the scientist and mission design engineer work rou-

tinely with data, theoretical models, and orbits (including orbital parameters, payload configuration, instrument fields-of-view, coordinated ground stations, etc.). At any given time they may work exclusively with models, data, or orbits, or combinations of all three. The most likely scenario involves an interactive combination of the three filtered through a mathematical interpreter and displayed through a broad set of visualization tools. This end-to-end approach of SAVS is illustrated in Figure 1, where the products of the SAVS system serve a range of NASA needs including: (1) data synthesis; (2) mission planning; (3) science and engineering trade studies; (4) data-model comparisons; and (5) model test and validation.

In mimicking the thought processes of the NASA scientist, the SAVS architecture has three entry concourses for daily activities: "Data," "Models" and "Orbits." By a "push of a button" in the SAVS front end scientists may enter any one of the three concourses and begin a procedure involving SAVS-guided functionality, whereby they can exercise an unlimited number of loops in any given concourse. With this approach, the scientists can treat and overlay multiple model runs or data sets, and visualize the results in any number of formats.

This SAVS concourse methodology and its generic products are illustrated in Figure 1.

The SAVS concept is designed for broad appeal and functionality since its architecture accommodates a general scientific approach to data access, manipulation, analysis, and visualization. Every scientific discipline deals with data, either in scalar, vector, or image format. Every discipline needs to develop, test, and apply empirical and first-principle model descriptions of cause-effect relationships. Every science needs to deal with the spatial and temporal coordinates of data sets (generally represented by a satellite's orbital elements, payload configuration, and instrument fields-of-view in a space science application); and every scientific discipline requires a suite of mathematical tools to handle data and prepare it for visualization. This is the generic system capability of SAVS, making it applicable to any scientific discipline. In the application described here SAVS is tailored to the needs of the space physics community with user interests in data from past programs (e.g., DE, IMP-8, ISEE, CRRES, etc.), in data provided by recent and imminent satellite operations (e.g., Wind, Polar, Geotail, etc.), and in effective planning of future missions (e.g., TIMED, IMI, and Grand Tour Cluster).

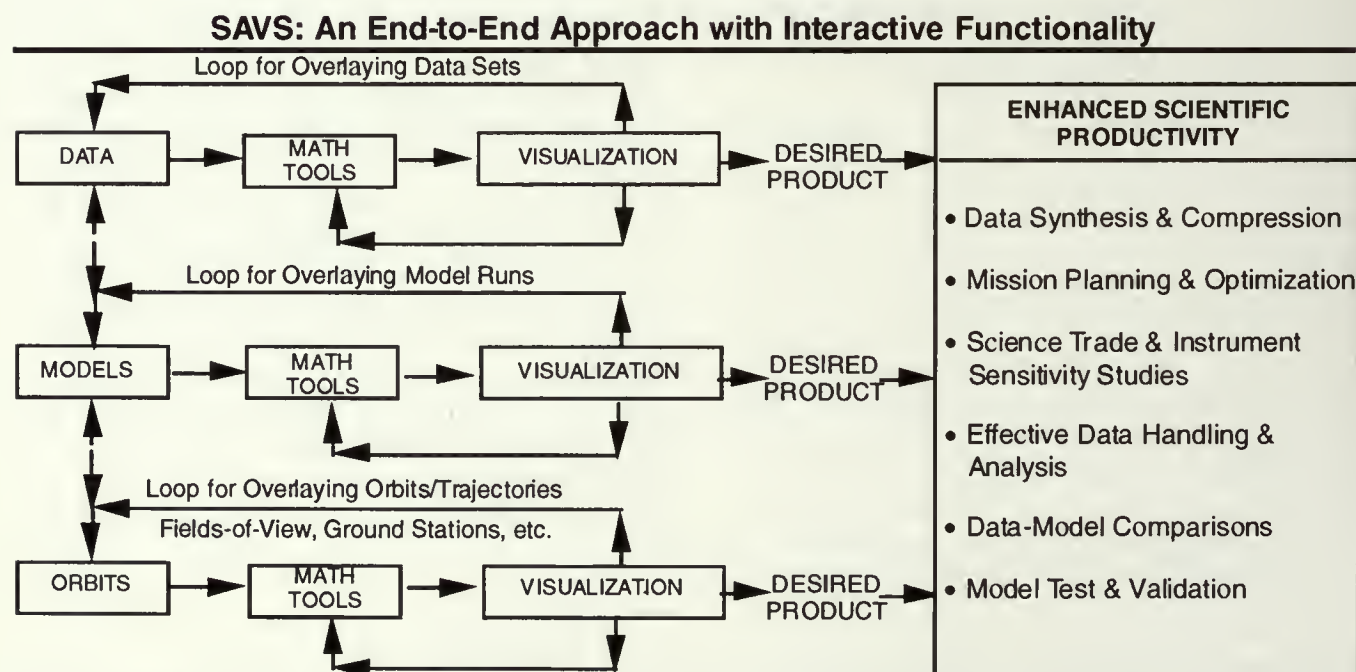


Figure 1. The SAVS entry concourses of "Data," "Models," and "Orbits" allow focus while providing flexibility for user-defined overlays of data, model, and experiment parameters in any number of visualization formats. The interactive functionality provides an end-to-end approach to enhanced scientific productivity.

1.3 The Visualization Software

Underlying data handling and analysis in the SAVS system is the AVS visualization software. Designed for a distributed network environment, AVS supports visualization applications running on a single system or across a network of systems. It provides a complete image display capability, including real-time pan and zoom, rotation and transformation, flip-book animation, and support for 8-bit, 24-bit, and floating point images. Imaging filters include operations, such as contrast stretching, pseudo-coloring, and histogram balancing, as well as data resizing operations, such as interpolation, cropping, and sampling. AVS provides tools for rendering volume data, a real-time isosurface generator, a unique transparent volume renderer which creates real-time semi-transparent images with full rotational and lighting control, and generation of geometric objects such as arbitrary slicing surfaces, dot surfaces, and vector nets. The software is based on industry-standard graphics, windowing, and programming interfaces including X-Windows and PHIGS+. (For details contact P. Esdale at AVS, Inc., via paule@avs.com.)

With AVS as the pivotal visualization tool, SAVS and its users also benefit from the extensive AVS module library at the International AVS Center (IAC) at the North Carolina Supercomputer Center (see NCSC/IAC in references), which serves as a catalyst for increasing AVS product functionality. As a worldwide repository available at no cost to AVS/SAVS users, the center collects, sorts, and distributes user-contributed public-domain modules. An example of existing aspects of this resource is the availability of 229 Khoros-derived AVS modules (Myerson and Reid, 1993). Khoros, like AVS, is a data-flow-oriented application, but it makes available to AVS and SAVS its one-dimensional and two-dimensional signal processing algorithms, complementing the well-established three-dimensional volume visualization capabilities in AVS.

2. THE "PUSH-BUTTON" FRONT END ARCHITECTURE

The background plate in Figure 2 shows portions of the opening SAVS panels. The top-

center of the *main panel* provides a "push-button" entry to any of the three SAVS concourses (Data, Models, and Orbits). The main panel also includes the display window (which is hidden by the front panel) where objects and object overlays are presented and where discrete controls can be used to manipulate the scene with definable increments of rotation, translation, magnification, and reduction.

The *pull-down menu bar* (at the top of the main panel ... "controls," "filters," etc.) provides access to a wide spectrum of functions: everything from general application needs to math and visualization tools for multi-variant/multi-dimensional data sets.

The *applications panel* (vertical panel on left) provides all parameter controls for each model run, for math tools used, and for the visualization process exercised. Behind these controls, and most SAVS buttons, are tailored networks transparent to the user, enabling him to easily assemble the necessary components for problem analysis without knowledge of the AVS-network building requirements faced by non-SAVS users.

The front panel in Figure 1 illustrates the staging process in a user's entry and exercise of the "Models" concourse. Selection of "Models" presents a choice browser at the top-left of the display window which allows the user to search through and select any number of models, including: (1) those in the SAVS library, (2) those which are user-unique and implemented within SAVS by built-in hooks, or (3) those accessed through remote procedure calls. In the current SAVS library are: (1) the Feldstein auroral oval model (Holtzworth and Meng, 1975), (2) the International Geomagnetic Reference Field (IGRF, Peddie, 1982), (3) the International Reference Ionosphere (IRI, Rawer and Ramanamurty, 1985), (4) the MSIS model representation of thermospheric densities and composition (Hedin, 1983), (5) the WIND model of thermospheric winds (Hedin, 1991), and (6) files of MHD model runs of the magnetosphere (Fedder and Lyon, 1987).

Selecting WIND (as an example of a SAVS application scenario) presents the user with: (1) an automatic visualization default selection of input parameters and a corresponding default visualization format (see pull-down radio buttons under the "Models" concourse button in the front




Figure 2. The opening SAVS panel (background plate) gives the user access to the “Data,” “Models,” and “Orbits” concourses, along with push-button utilities for data and model handling and visualization techniques. The front plate is an example of a SAVS application scenario discussed in the text.

panel of Figure 2), (2) the default object in the display window, and (3) the control panel on the left which allows the user to tune the model and its image according to all its driving variables. For WIND, the input variables include dial controls for universal time (UT), the activity index (Ap), and the 10.7 cm solar flux (F10.7). The control panel also offers keyboard controls for selection of computational resolution, altitude for a surface display of the thermospheric wind vectors, day-of-the-year, and scaling of N-S and E-W components of the velocity vectors. The user can pan, zoom, or rotate the image, exercise a parameter study by

“instantaneously” changing the input conditions and observing the results, or begin to overlay data or other models (e.g., a color rendering of O_2 and N_2 densities from MSIS, the boundaries of the auroral oval according to Feldstein (with controls over UT and Kp), the IGRF to study meridionally- or zonally-coupled regions, or data for model-measurement comparisons). An illustration of multiple overlays will be taken up in the discussion of Figure 4.

In the case of model-measurement comparisons, SAVS has a specialized module that projects the results of any three-dimensional model onto any

one-dimensional line (Szuszczewicz, et al., 1993), a requirement for comparisons of along-track satellite observations with large-scale three-dimensional model predictions. Under the “Orbits” concurrence, the user can call up any predefined mission orbit or one defined by the theoretical CADRE-3 satellite ephemeris package (Galperin, 1990) imbedded in the SAVS “Orbits” concurrence. The user can then implement the three-dimensional to one-dimensional interpolation module by activating the “three-dimensional to one-dimensional” item in the “filters” submenu of the menu bar.

3. DATA ACCESS

The pivotal point of any analysis effort involves the data, with the SAVS architecture making available the tools to browse, access, and read data from local and remote repositories. While data being accessed by NASA researchers exist in an exceedingly wide range of formats and organizational levels, HDF (hierarchical data format), CDF (common data format), and netCDF are developing broad appeal. HDF, for example, is the baseline standard for EOSDIS data product generation, archival activities, ingest, and distributions. And CDF has been adapted by ISTP for many of its data products. With this perspective, we are providing HDF, CDF and netCDF readers, while including a set of “hooks” so that users can include readers for “non-standard” data formats.

To access data from national and international sites, we are integrating browsing and remote procedure calls into SAVS via adaptations of a menu-oriented data base front-end. Users, with widely varying applications, can therefore act as their own data retrieval service, thus relieving a longstanding bottleneck connected with poor user access to the data and dependence on data centers for satisfying their requests. While this element of SAVS is still under development, a prototype system has been implemented based on the relational data base management techniques employed in SAIC’s CenterView (Brennan, 1987). CenterView has a history of more than ten years of development efforts and is routinely used as the hub of a worldwide network which acquires, stores, and analyzes an active real-time global seismic data base.

We illustrate the ease and utility of this SAVS element of remote data access, which we call SAVS-View, in an application that can be thought of as generic to interplanetary and magnetospheric physics. Working with the National Space Science Data Center (NSSDC) at the NASA Goddard Space Flight Center we set up a remote library of data and developed an index system and a corresponding set of “keys” which would allow a user to browse the index catalogs and retrieve the desired subset of data from the library. This browse-and-retrieve activity involved a number of successful tests and demonstrations, one of which employed an IBM RS6000 at the University of Colorado linked with an Oracle-based network server at NSSDC.

The background panel in the upper plate of Figure 3 shows the set of keys that our “test user” activated in establishing a query of available data in our library. The panel in the foreground shows the resulting list of data which matched his index of “keys” after having launched the query. In this case, the user was interested in solar wind input to the magnetosphere and wanted to build an empirical picture of that energy using available data and the solar wind energy formalism represented by the epsilon parameter (Perrault and Akasofu, 1978; Vasyliunas et al., 1982) which has the form $\epsilon = 1/2 v B^2 \sin^4(\Theta/2)$. In this relationship v is the solar wind velocity, B is the value of the interplanetary magnetic field, Θ is the polar angle of the component of the IMF normal to the Sun-Earth line and measured from the northward geomagnetic axis, and l_0 is a constant. In order of selection, the user struck the “data” key, and then the “subcategory” key which allowed his choice or choices of measurements. Here, he selected the total B-field, the solar wind velocity and density, and the B_z -component of the IMF as an indicator for Θ . His data source of interest was IMP-8, so he keyed “Missions” and selected IMP-8, and then launched the query using the pull-down query menu at the top of the panel. The entire effort, including the time for SAVS-View to find and list the available data, required approximately two minutes. The transfer of those files and the presentation of the B and B_z parameters using the SAVS two-dimensional stack plot capabilities are presented in the lower plate in Figure 3. The formatting of the data plots took advantage of SAVS tools that

allowed the selection of any parameter in the file for the abscissa and the ordinate. SAVS also allows truncation of the individual plots, as well as their extensions along any user-selected domain of the independent variable.

In the development of data access tools, current SAVS activities also include the implementation of

a SAVS utility to access the Logical Library System (LLS) developed at the NASA Goddard Space Flight Center (Jacobs, 1993). The LLS, as the name suggests, has the functionality of a library, with programs, missions, events, etc. catalogued within the framework of library shelves (e.g., magnetospheric physics), books (e.g., the DE and

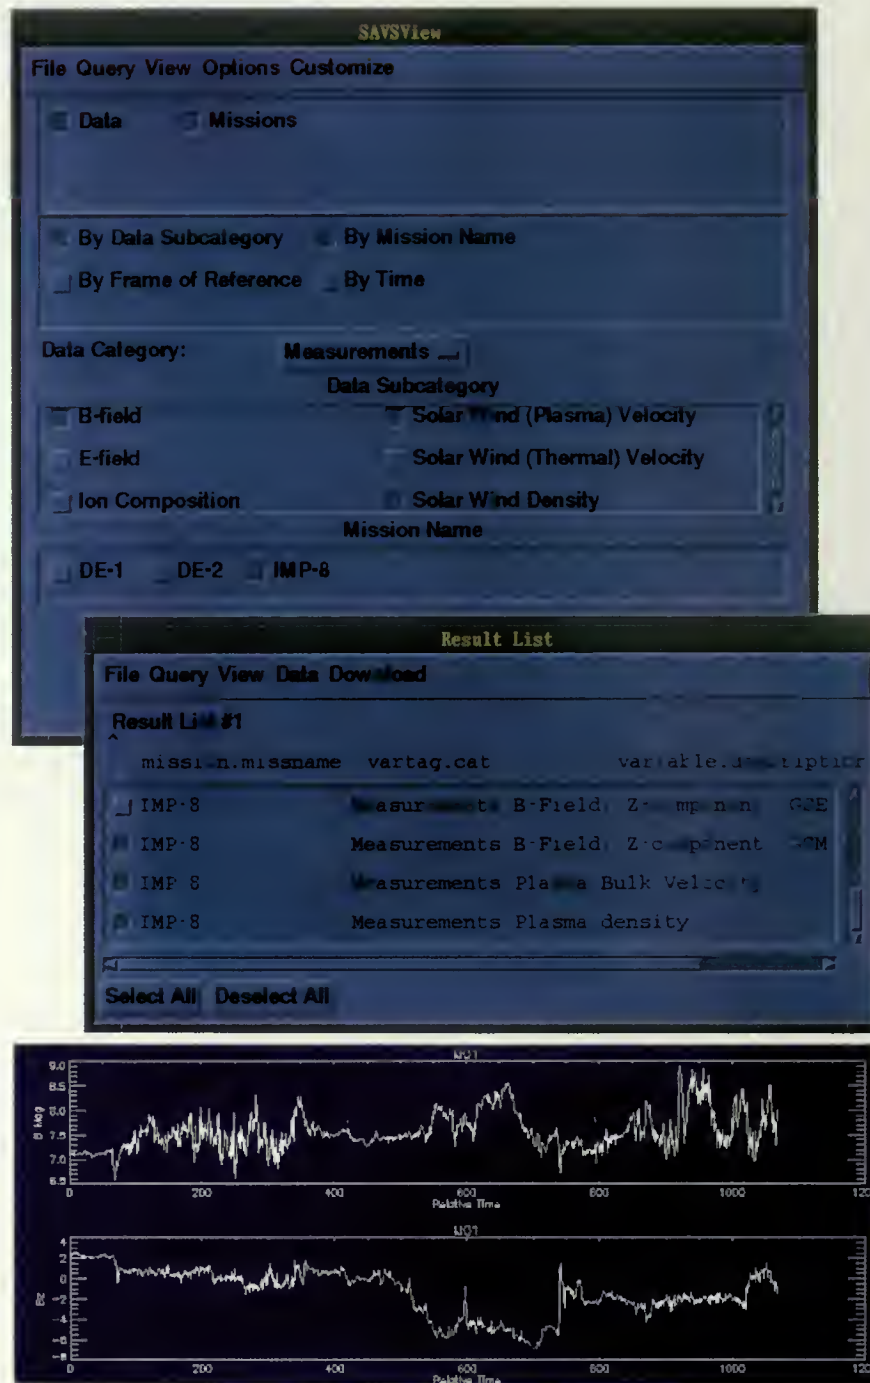


Figure 3. The upper plate shows the set of "keys" that a test user activated in establishing a query of available data at a remote site. The bottom panel shows a subset of the data acquired, manipulated and displayed within the SAVS system.

IMP-8 missions), chapters (e.g., mission years or event categories), and pages (e.g., specific magnetic storms). This system takes advantage of the GSFC/NSSDC Master Directory and will eventually make data at the center available online to SAVS calls.

4. VISUALIZATION APPLICATIONS TO MISSION SCENARIOS

We demonstrate the interactivity and extensibility of the SAVS system with an application scenario that can be considered generic to any mission planning exercise or to an activity involving data reduction and analysis that requires: (1) a number of models, (2) ground-based and spaceborne data, (3) “in situ” and remote sensing techniques, (4) perspectives on fluxtube coupled domains, and (5) co-registration of satellite-borne and ground-based diagnostics.

A strawman scenario for a given mission is presented in Figure 4. The scenario has counterparts in the CRRES, TIMED, or ISTP missions, requiring the coordination of satellite passes with ground-based optical, ionosonde, and radar diagnostics. The “Models” concourse was used in generating the object in the left panel in Figure 4 by overlaying: (1) thermospheric winds (blue vectors) and oxygen densities (global color-coded surface) at 350 km using the SAVS library models WIND and MSIS, (2) the auroral oval boundaries from the Feldstein

model, (3) altitude cut-planes of ionospheric densities from the IRI in the latitude and longitude planes intersecting the coordinates of the Arecibo Observatory, and (4) geomagnetic field lines (in white) from the IGRF. The satellite trajectory was generated in the “Orbits” concourse using the CADRE-3 package.

With push-button and dial controls the user can: (1) adjust the auroral oval for any magnetic activity index or universal time, (2) change the heights of the thermospheric wind and density calculations, (3) vary the activity index A_p , the solar F10.7 flux, and the day-of-the-year input drivers for WIND and MSIS, (4) alter the year for the IGRF, (5) tune the month and sunspot number for the IRI, and (6) change the coordinates for the intersecting cut-planes (with interest, for example, in high-latitude passes in conjunction with Sondrestrom radar operations in Greenland). The cut-planes could also be in magnetic E-W and N-S directions instead of being aligned with the geographic contours of latitude and longitude at the Arecibo Observatory.

From the CADRE-3 orbit specifications the user can develop stack plots of any number of relevant parameters versus UT (or for example, versus latitude and longitude, or versus MET, or versus LT). Parameters available to the user include: (1) SZA, MLAT, MLONG, MLT, and L-shell value, (2) point of conjugacy, (3) spacecraft velocity vector relative to the ambient B-field, (4) height of the





Figure 4. A visualization application of SAVS to a mission scenario with counterparts in the CRRES, TIMED, and ISTP missions. The left panel provides a global view, while the right panel focuses on the conjunction between the satellite pass and the radar field-of-view at the Arecibo Observatory.

terminator relative to the satellite track, (5) B-field values or values of its E-W and N-S components, etc.

The user can also call-up the three-dimensional to one-dimensional module to interpolate any and all model results along any segment of the satellite track and through the “Data” concurrence compare the results with local or remote data bases of along-track measurements.

For co-registration with ground-based diagnostics, the right panel of Figure 4 presents a zoomed-in perspective on the Arecibo site that includes the projection of the ground-based HF radar beam up to an altitude of 400 km. That projection is developed and displayed using the SAVS “field-of-view” module that is controllable by the user to represent any and all ground-based remote sensing diagnostics. The user need only input the coordinates of the ground site, the azimuth and elevation of the instrument’s line-of-sight, and its beam-width (or acceptance) half-angles in the N-S and E-W directions.

Inspection of the right panel of Figure 4 shows that the satellite (the black sphere) at the given UT in its trajectory is not on a fieldline that connects to the region diagnosed by the ground-based radar. With animation and pause controls, the user can determine the exact time of fieldline-coupling and begin a number of SAVS-controlled operations. For example, the three-dimensional to one-dimensional module can be used to interpolate ionospheric and thermospheric model densities, temperatures, and composition onto the connecting fieldline, and the model data used as inputs exercised to calculate fluxtube-integrated conductivities from the satellite to the region covered by the ground-based diagnostics. The module could also be used to calculate integrated densities or emissions along the field-of-view of the ground-based sensor. The field-of-view module can also be “attached” to the spacecraft to project and visualize the field-of view of an onboard remote sensing device. In this case, the animation and visualization product could focus on the determination of the time of intersection of the fields-of-view of the satellite and ground-based sensors. The scenarios are unlimited. Models can be tuned and retuned, and results visualized in simple stack plots of along-track correlations for comparisons with data. Visualizations can include constant

altitude surfaces of any geophysical parameters (e.g., N_e , T_e , and densities of O^+ , N_2^+ , ..., N_n , T_n , and densities of O , N_2 , ..., thermospheric wind vectors and their meridional and zonal components, etc.). Instead of constant altitude representations, the user can choose altitude cut-planes, isosurfaces, or isocontours. The proper visualization product is only determined by the perceptive and inquisitive scientific mind, with the freedom of the push-button operations of the SAVS system making the choices immediately available.

5. COMMENTS AND CONCLUSIONS

As described, the SAVS architecture accommodates a generic scientific approach to data access, analysis, and visualization with tools that provide online mathematical analysis (e.g., averaging, interpolating, spectral analyses through FFTs, etc.) along with one-, two-, and three-dimensional displays that can include animation. Data displays and model results can be rendered in an array of two-dimensional stack plots with the user exercising control of abscissa and ordinate selection from the full set of parameters in the data file. The observations and the model results can be rendered in color codes, isolines, isosurfaces, cut-planes, and texture maps, with controlled lighting and transparency. The three-dimensional to one-dimensional interpolation modules can be used for along-track model-measurement comparisons, field-line-integrated calculations, or intensity calculations along the line-of-sight of remote sensing devices. And the field-of-view module can be used to establish regions of co-registration involving multi-satellite missions and ground-based observatories.

The current system has been developed and tested using a relatively large number of models covering magnetospheric, ionospheric, and thermospheric physics (see e.g., Szuszczewicz et al, 1993 and Szuszczewicz, 1994). And tests involving the data readers and local and remote data access tools have included the CRRES, IMP-8, DE, and the ISEE missions.

The push-button architecture is intended to service user needs focused on science, and not graphics network building and programming. The complexities of the tools are in the background as specialized and tailored modules within the SAVS

system framework. The “expert user,” however, can still access the foundations of AVS and build his or her own graphics network.

While SAVS includes a library of models and data readers, it is also providing the hooks and menu driven recipes for user-unique formats and models in order to guarantee flexibility and extensibility in the overall system. The integrated system is generic in its capabilities with natural applications to ISTP, TIMED, EOS, and TRMM, to name only a few of the near-term large-scale science programs at NASA.

Acknowledgments. This work was supported by NASA’s Applied Information Systems Research Program (AISRP) under contract NAS5-32337.

REFERENCES

- Brennan, M.A., “Center for Seismic Studies Database Structure Version 2.8,” Tech. Rpt. C87-04, Teledyne-Geotech, Arlington, VA (September 1987).
- Fedder, J.A. and J.G. Lyon, “The Solar Wind - Magnetosphere - Ionosphere Current - Voltage Relationship,” *Geophys. Res. L.* 14, 880 (1987).
- Galperin, Y. “An IBM PC - Compatible Software Package for Satellite Orbit Definition Ephemeris Determination, and Along-Track Geophysical Parameters,” (private communication, 1990).
- Hedin, A.E., M.A. Biondi, R.G. Burnside, G. Hernandez, R.M. Johnson, T.L. Killeen, C. Mazaudier, J. W. Meriwether, J.E. Salah, R.J. Sica, R.W. Smith, N.W. Spencer, V.B. Wickwar, and T.S. Virdi “Revised Global Model of Thermosphere Winds Using Satellite and Ground-Based Observations,” *J. Geophys. Res.* 96, 7657-7588 (1991).
- Hedin, A.E., “A Revised Thermospheric Model Based on Mass Spectrometer and Incoherent Scatter Data: MSIS-83,” *J. Geophys. Res.* 88, 10710 (1983).
- Holtzworth, R.H. and C.I. Meng, “Mathematical Representation of the Auroral Oval,” *Geophys. Res. Lett.* 2, 337 (1975).
- Jacobs, B.E., “Implementation Plan for the Space Science Data Digital Library (SSD of lib),” NASA/GSFC Draft Report (November 16, 1993).
- Myerson, T. and D. Reid, “New Khoros Modules for AVS,” *AVS Network News* 1, Is. 3, 6-10 (1993); distributed by the University of Manchester through the anonymous FTP address ftp.mcc.ac.uk.
- NCSC/IAC, North Carolina Supercomputing Center, Int’l AVS Ctr., 3021 Cornwallis Rd., Research Triangle Park, NC 27709, anonymous FTP to avs.ncsc.org.
- Peddie, N.W., “International Geomagnetic Reference Field: the Third Generation,” *J. Geomag. Geoelectr.* 34, 309-326 (1982); for revisions contact the Nat’l. Geophys. Data Ctr., Boulder, CO (303-497-6478).
- Perreault, P. and S.-I. Akasofu, “A Study of Geomagnetic Storms,” *Geophys. J. Roy. Astron. Soc.* 54, 547 (1978).
- Rawer, K. and Y.V. Ramanamurty, “International Reference Ionosphere - Status 1985/86,” URSI/COSPAR Workshop Proceedings, Belgium, 28 October 1985, Wheaton & Co., Ltd., UK, 1985.
- Szuszczewicz, E.P., A. Mankofsky, P. Blanchard, C. Goodrich, D. McNabb, and D. Kamins, “SAVS: Progress and Plans,” in AISRP Workshop III Proceedings, G. Mucklow, editor, Published by NASA Information Systems Branch, Washington, DC (1993).
- Szuszczewicz, E.P., “Visualization: A Pathway to Enhanced Scientific Productivity in the Expanding Missions of Space and Earth Sciences,” in *Visualization Techniques in Space and Atmospheric Sciences*, E.P. Szuszczewicz and J. Bredekamp editors, NASA Printing Office, (Washington, DC, 1994).
- Upson, C., T. Faulhaber, Jr., D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A.V. Dam, “The Application Visualization System: A Computational Environment for Scientific Visualization,” *IEEE Computer Graphics and Applications*, 30-42 (July 1989).
- Vasyliunas, V.M., J.R. Kan, G.L. Siscoe, and S.-I. Akasofu, “Scaling Relationships Governing Magnetospheric Energy Transport,” *Space Sci.* 30, 359-365 (1982).

Visualization of International Solar-Terrestrial Physics Program (ISTP) Data

RAMONA L. KESSEL and ROBERT M. CANDEY
*NASA Goddard Space Flight Center, Code 632,
 Greenbelt, MD*

SYAU-YUN W. HSIEH and SUSAN KAYSER
Hughes STX, Greenbelt, MD

The International Solar-Terrestrial Physics Program (ISTP) is a multi-spacecraft, multi-national program whose objective is to promote further understanding of the Earth's complex plasma environment. Extensive data sharing and data analysis will be needed to ensure the success of the overall ISTP program. For this reason, there has been a special emphasis on data standards throughout ISTP. One of the key tools will be the Common Data Format (CDF), developed, maintained, and evolved at the NSSDC, with the set of ISTP implementation guidelines specially designed for space physics data sets by the Space Physics Data Facility (associated with the NSSDC). The ISTP guidelines were developed to facilitate searching, plotting, merging, and subsetting of data sets. We focus here on the plotting application. A prototype software package was developed to plot Key Parameter (KP) data from the ISTP program at the Science Planning and Operations Facility (SPOF). The ISTP Key Parameter Visualization Tool is based on IDL and is keyed to the ISTP guidelines, reading data stored in CDF. With the combination of CDF, the ISTP guidelines, and the visualization software, we can look forward to easier and more effective data sharing and use among ISTP scientists.

1. INTRODUCTION

The International Solar-Terrestrial Physics Program (ISTP) is a multi-spacecraft, multi-national program whose objective is to promote further understanding of the Earth's complex plasma environment. The first ISTP spacecraft, the Japanese-U.S. (ISAS-NASA) Geotail spacecraft, is now in orbit. Two NASA spacecraft, WIND and POLAR, are slated for 1994 launches. The fleet of spacecraft also includes several equatorial spacecraft: (NOAA) GOES spacecraft and (DOE) LANL spacecraft. IMP-8 supplements these observations. Ground-based data from DARN, SESAME, CANO-PUS, and Sondstromfjord, and theoretical studies and modeling complement the spacecraft data. In addition, NASA is collaborating with the European Space Agency (ESA) in two additional missions: SOHO and Cluster. The Russian Space Research Institute (IKI) is participating through the Inter-Agency Consultative Group with the Interball mission. All of these missions together will study the generation, flow, and dissipation of mass, momentum, and energy between the Sun and the Earth.

Extensive data sharing and data analysis will be needed to ensure the success of the overall ISTP program. For this reason, there has been a special emphasis on data standards throughout ISTP. One of the key tools will be the Common Data Format (CDF), with the set of ISTP implementation guidelines specially developed for space physics data sets by the Space Physics Data Facility (associated with the NSSDC). CDF (along with the ISTP guidelines) is a means of storing data and metadata (description of the data) in a standard, readily accessible format. CDF, developed, maintained, and evolved at the National Space Science Data Center (NSSDC), is a

management tool for scientific data sets. Data are organized into an interlocking grid structure with standard labels, time tags, dependencies, uncertainties, and offsets explicitly defined for later searching, plotting, merging, and subsetting. CDF has a high-level programming and toolkit library which provides functionality without excessive programming and is portable to a wide variety of platforms. CDF data sets are transportable between any of the platforms that CDF supports (VAX, Sun, SGi, IBM, HP, and PC). CDF is freely available to the public, and will be supported at least throughout the lifetime of ISTP and the IACG campaigns.

Many tools now exist or will be designed and implemented in the future for use with CDF. Additionally, there are specific tools designed and planned for CDF data sets which include the standard set of ISTP descriptions. All of these are shown in Figure 1. This figure is divided into three sections: creation, visualization, and manipulation.

The C and Fortran programming interface spans all of these. The ISTP Guidelines add a layer of specific metadata (attributes) to the CDF data sets; the guidelines are shown in blue next to the C and Fortran programming interface in Figure 1. The tools in the three sections use either generic CDFs or CDFs with ISTP Guidelines. We briefly describe each of these tools. The CDF toolkit, which spans all three sections, is part of the generic CDF distribution and includes routines for assisting in the creation of CDFs through the use of skeleton tables (section 2), tools to list and browse the data values and metadata (visualization), and tools to edit and convert CDFs (manipulation). These tools work well with any CDF. CDF data sets can be created just using the C or Fortran programming interface, although the process is simplified by using the combination of skeleton tables (CDF toolkit) and programming. The creation process is made even easier through a prototype tool called Make_CDF,

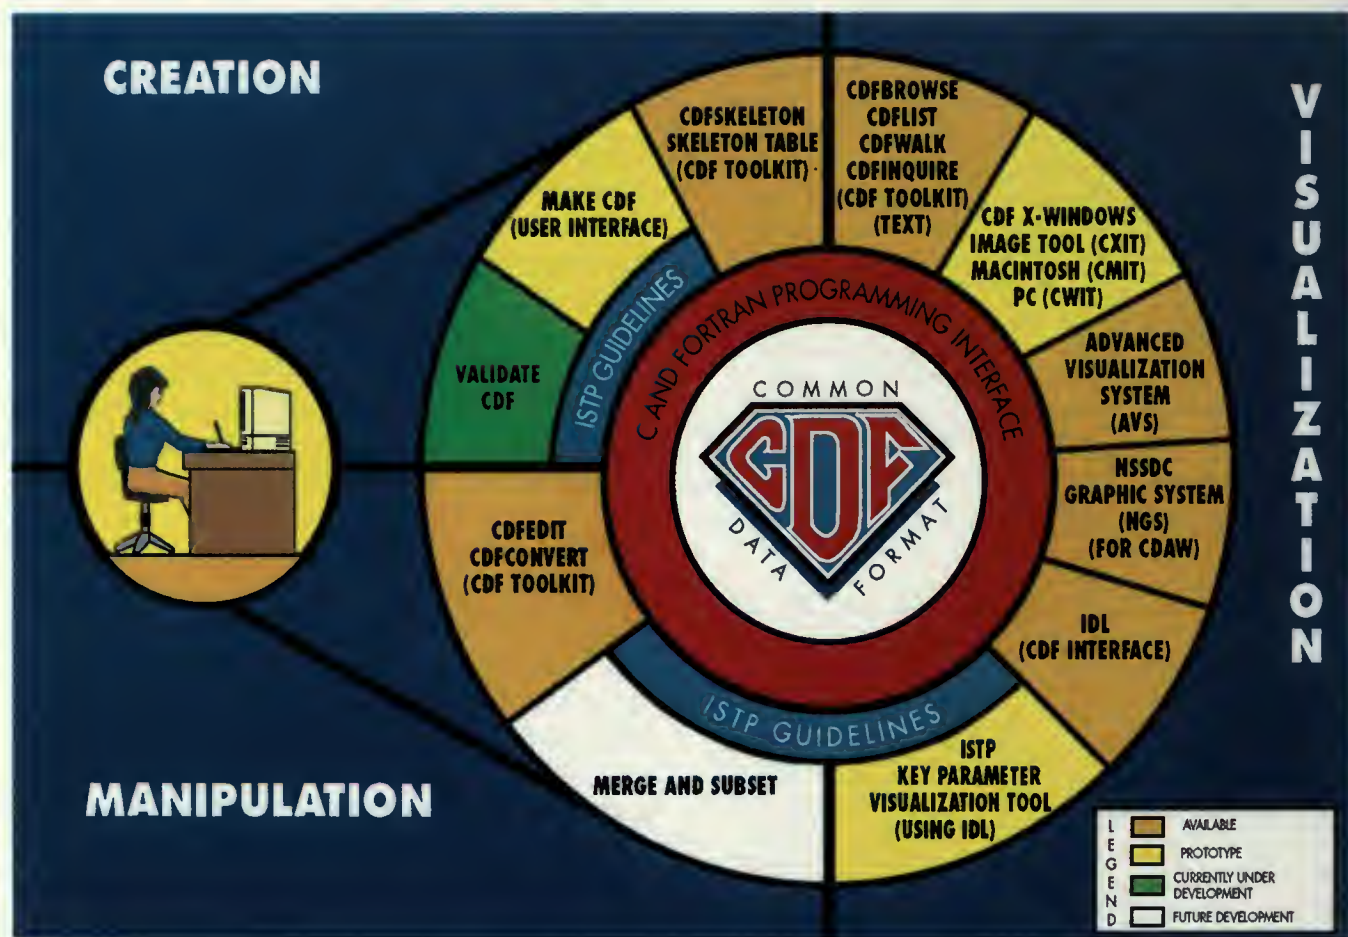


Figure 1. CDF associated tools and programs: white background indicates tools in the planning stage, light colored background indicates software development in progress, and somewhat darker background indicates prototype software.

which also uses the ISTP Guidelines (shown in yellow in Figure 1). With a data set in hand, a user describes the input and output format through a user interface, and the tool then creates the CDF data set. No programming is required. Other tools built on top of CDF and the ISTP Guidelines include a package which is being developed to validate whether or not a particular CDF meets the ISTP guidelines (i.e., whether or not all of the required metadata is included), a Merge and Subset package which is in the planning stage, and the ISTP Key Parameter Visualization Tool described below. This last tool presently plots time series key parameter data via the Interactive Data Language (IDL). There are other packages which make use of the CDF library, but do not take advantage of the metadata included in files designed with the ISTP Guidelines. For example, three visualization tools which make use of CDF are shown in Figure 1: CDF X-Windows Image Tool (CXIT, available free from the NSSDC) for viewing image data; NSSDC Graphics System (NGS, available online at NSSDC) for plotting data from the Coordinated Data Analysis Workshop (CDAW), and the Advanced Visualization System (AVS) which is a commercial package for plotting multi-dimensional data. In addition, IDL has a built-in CDF interface, as well as access to the programming library.

The ISTP Key Parameter Visualization Tool is a prototype software package (version 1.1) developed at the Science Planning and Operations Facility (SPOF), for plotting ~1 minute resolution survey data from the ISTP program (referred to as Key Parameter (KP) data). This tool consists of two parts: one written in Fortran that is responsible for reading the data and metadata; and one written in IDL with widgets that provides the data processing and plotting capabilities. In addition, this tool has its own custom interface between IDL and the CDF software, and does not use either the CDF library provided by IDL version 3.0 nor the IDL interface programs provided by CDF version 2.3. These versions were not available when the tool was first designed, and then it was decided to preserve the independence of the custom interface to allow for use with PV-Wave or other commercial plotting packages. The data files must be in CDF and compliant with the ISTP standard in order to be plotted by the visualization tool. The tool uses the

metadata as outlined in the ISTP guidelines: labels, units, scales, etc. The visualization tool is available to ISTP investigators through the Central Data Handling Facility (CDHF) on an as-is basis with no support. The visualization tool was developed on a DEC workstation running Ultrix and later ported to a Sun Workstation. On both platforms, IDL with an X-Windows environment, is required.

At this time, and for some time in the future, the ISTP data sets are proprietary. It is envisioned that the non-proprietary CDAW-8 and CDAW-9 data which were originally cast in CDF will be updated to embrace the ISTP guidelines and can then take full advantage of the visualization tool. Other investigators will be able to take advantage of these tools by putting data into CDF with the ISTP guidelines. We provide in this paper a discussion of how to create ISTP-like CDFs, and how to visualize the data. The plots are either time series data or hodogram type plots. We show a sample plot and discuss the planned extensions.

2. CREATION OF ISTP-LIKE CDFs

ISTP CDF data sets are organized as daily files, with separate files for each experiment on each spacecraft and for each ground-based station. The data are, for the most part, one minute averaged time series scalar and vector data, although some images at a resolution of approximately 5 minutes are also being provided. For the new spacecraft within ISTP, the orbit and attitude files are being provided separately (also in CDF), whereas the previously launched spacecraft (Imp-8, GOES, LANL) provide the orbit information along with their data variables. Each ISTP CDF data set includes both data and descriptions of the data (metadata). The descriptions include header information (e.g., number of variables, number of attributes, dimensions, data encoding), global information about the data set (and experiment) as a whole (global attributes) and information carried with each variable (variable attributes). For ISTP, there is a standard set of attributes to promote easier and more effective data sharing, referred to as the ISTP guidelines. Users expect and will receive the information required to both understand and analyze each data set. In a short paper we cannot provide all of the details necessary to create a CDF

data set adhering to the ISTP guidelines. Instead, we give the flavor of the creation process through analogy. The Make-CDF interface, mentioned previously, eliminates the coding step discussed below. We also list the complete set of ISTP standard attributes associated with a scalar variable. Some of these attributes are not related to plotting, and some have not been incorporated into the present release of the visualization software.

2.1 CDF Analogy

The process of creating a CDF data set is similar to the process of building a house (see Figure 2). *The number of floors in a house is similar to the dimensions of the CDF, and the rooms are analogous to variables. Furniture and possessions are analogous to data.* For the ISTP

CDF designer, the data itself is always crucial to the design of the CDF. The data and the ISTP Guidelines together dictate the design. The global features of a house (such as fireplaces, decks, garage) are similar to the global attributes that define the data set as a whole. Particular room features (closets, outlets, windows, sinks, bathtubs) are analogous to the variable attributes; which attributes to include depends on the nature of the variable. Just as the blueprints are the plans for a house, the skeleton table is the plan for the particular CDF data set.

Once the blueprint is finished, the building stage can begin. This is usually the most time consuming part of the process for both houses and CDFs. For CDFs, the building stage involves changing the ASCII skeleton table description into a binary version and writing the computer code to add the data to the description. The final stage

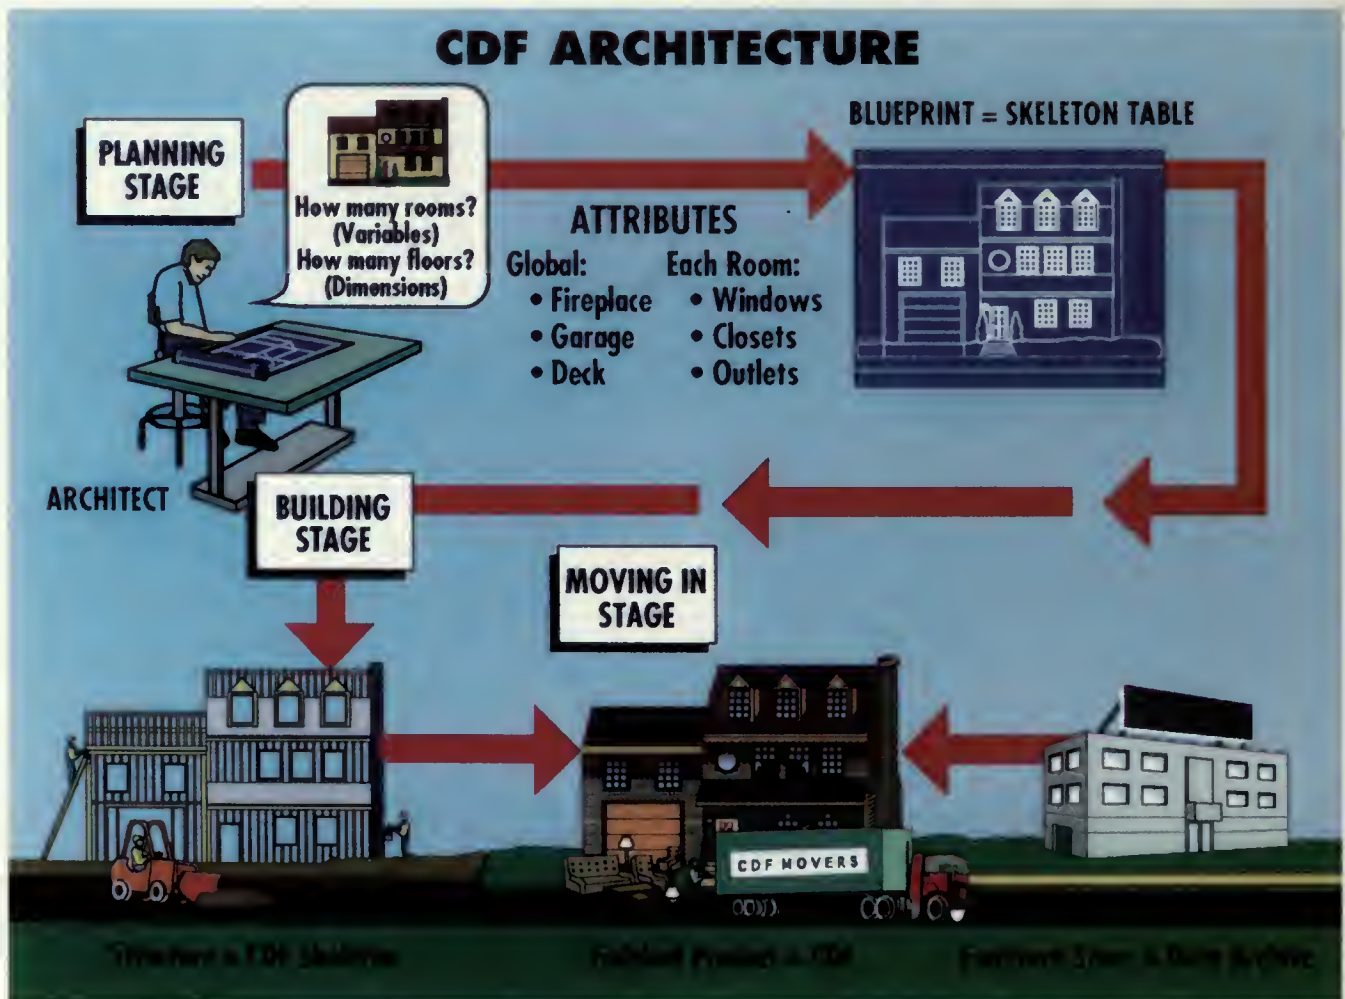


Figure 2. Creating a CDF data set in analogy to building a house.

involves moving furniture into the house or data into the CDF (running the computer code). Once the furniture has been moved in, the house is ready for occupancy; when the data is put into the CDF, the CDF data set is complete and can be viewed, ported to another location, manipulated, or visualized via plotting software.

2.2 Standard Attributes

Variable-scope-attributes are linked with each individual variable, and are an important means of providing additional information about each variable. A standard set of these attributes (required for every variable) is very important, for this is where certain required information can be stored in a commonly defined manner. Any application program can utilize the information, but the CDF data set designer does not need to know the details of the application program. The specific variable-scope-attributes in the ISTP standard list for scalar data are listed in Table 1 and defined below.

FIELDNAM (required) holds a character string (up to 30 characters) which is longer and more

descriptive than the name of the variable. It can be used to label a plot either above or below the axis, or can be used as a table heading. Therefore, consideration should be given to the use of upper and lower case letters where the appearance of the output plot or table heading will be affected.

LABLAXIS (required) should be a short character string (no more than 10 characters, but preferably 6 characters) which can be used to label a y-axis for a plot or to provide a heading for a table.

UNITS (required) is a character string (no more than 20 characters, but preferably 6 characters) representing the units of the variable, e.g., nT for magnetic field. If the standard abbreviation used is short then the units value can be added to a table heading or plot label.

VALIDMIN and VALIDMAX (required) hold values which are, respectively, the minimum and maximum values for a particular variable that are expected over the lifetime of the mission.

SCALEMIN and SCALEMAX (required) are values which can be based on the actual values of data found in the CDF data set or on the probable uses of the data, e.g., plotting multiple files at the same scale. The visualization software uses these attributes as defaults for plotting, but lets the user override the default scale.

FILLVAL (required) is the number inserted in the CDF in place of data values that are known to be bad or missing. Fill data are always non-valid data.

FORMAT (required) is the output format used when extracting data values out to a file or screen (using CDFlist). The magnitude and the number of significant figures needed should be carefully considered. A good check is to consider it with respect to the values of VALIDMIN and VALIDMAX attributes. The output format can be in Fortran or C, but Fortran is preferred.

DEPEND_0 (required for time-varying variables) explicitly ties a "parent" variable to the time variable on which it depends; the value of the attribute is the time variable. For ISTP data sets, the data is at one time resolution and the time variable in each data set is called "Epoch" (stored as milliseconds AD, but displayed as year, month, day, hour, minute, second).

VAR_TYPE (required) identifies a variable as either **data** (integer or real numbers) or as **metadata** (labels or character variables).

Table 1. ISTP Standard Set of Variable-Scope-Attributes for a Scalar Variable.

Scalar Variable Attributes for "Density"		
Attribute Name	Type	Example of Attribute Value
Required Attributes		
FIELDNAM	CDF_CHAR	Proton No. Density
LABLAXIS	CDF_CHAR	Np
UNITS	CDF_CHAR	no/cc
* VALIDMIN	CDF_REAL4	0.0
* VALIDMAX	CDF_REAL4	50.0
* SCALEMIN	CDF_REAL4	0.0
* SCALEMAX	CDF_REAL4	50.0
* FILLVAL	CDF_REAL4	-1.0E31
FORMAT	CDF_CHAR	F6.2
* DEPEND_0	CDF_CHAR	Epoch
VAR_TYPE	CDF_CHAR	data
DICT_KEY	CDF_CHAR	
Optional Attributes		
CATDESC	CDF_CHAR	Proton number density determined from a moment calculation
SCALETYP	CDF_CHAR	linear
AVG_TYPE	CDF_CHAR	standard
* DELTA_PLUS_VAR	CDF_CHAR	uncertainty
* DELTA_MINUS_VAR	CDF_CHAR	minus_uncertainty
* OFFSET_0	CDF_CHAR	time_offset

•Indicates that these are pointer class attributes which have as their value another variable in the CDF data set. These attributes are used to link the parent variable, in this case, Density, to the "attached variables": Epoch, uncertainty, minus_uncertainty, and time_offset. The * indicates that these attributes are of the same data type as the variable.

DICT_KEY (required) comes from a data dictionary and describes the variable to which it is attached. (At this time ISTP does not have a completed data dictionary; this attribute is included as a placeholder only and at this time does not affect visualization.)

CATDESC (optional) (catalog description) is an 80-character-max string which is a textual description of the variable. **This attribute should be used whenever there is possible confusion in the meaning of the variable or when FIELDNAM is not long enough to store the information.**

SCALETYP (optional) indicates whether the variable should have a **linear** or a **log** scale as a default. If this attribute is not present, **linear** scale is assumed.

AVG_TYPE (optional) sets up useful default conditions: different techniques appropriate to averaging different types of data. If this attribute is not present, **standard** average, i.e., simple arithmetic mean, is assumed. For other options, see [1].

DELTA_PLUS_VAR and DELTA_MINUS_VAR (optional) are included to point to a variable (or variables) which stores the uncertainty in (or range of) the original variable's value. The uncertainty (or range) is stored as a (+/-) on the value of the original variable. For many variables in ISTP, the original variable will be at the center of the interval, so that only one value (or one set of values) of uncertainty (or range) will need to be defined. In this case, DELTA_PLUS_VAR, and DELTA_MINUS_VAR will point to the same variable.

OFFSET_0 (optional) is used as a way to carry multiple time resolutions or multiple time tags offset from each other in a file, while maintaining only one time that is the record ordering parameter. The variable which holds the time offset(s) is the value of the attribute.

3. VISUALIZATION OF ISTP-LIKE CDFs


The visualization software is best described through a demonstration of its user interface. Figure 3 shows the IDL widget layout for the software. There is a main panel which consists of four sections labelled A through D in Figure 3. Section A is the structure setup which is a series of buttons with a number of options; these may be

selected in any order. Users choose the number of panels to plot and the type of plot: either a time series or a variable vs. variable plot. Users choose to plot either ISTP CDFs or Other CDFs. Users also choose one of three options: to plot variables from a single CDF (single CDF per panel per plot); to allow variables from multiple CDFs to be plotted, but only from a single CDF on any one panel (single CDFs per panel and multi CDFs per plot); or to plot from more than one CDF for any panel (multi CDFs per panel per plot).

Section B shown in Figure 3, is a series of selections each of which require another screen. The selections must be made in order from top down as shown in Figure 3, and listed below. A warning message will appear if one of the selections is skipped.

- a Select the number of key parameters per panel, up to two is allowed.
- b Select the spacecraft or ground-based investigation (ISTP mission).
- c Select the ISTP experiment.
- d Select the ISTP CDFs.
- e Select the ISTP key parameters (for time series plot).

For each of these selections, more choices are presented on subsequent screens; users are presented with choices so that they are not required to type in values which may not be known apriori. Two of these screens (c and d) are context sensitive, i.e., if IMP-8 is chosen as the spacecraft in b, then only two experiments appear in c: the magnetometer and the plasma experiment, because these are the only two instruments on IMP-8 presently supplying data to ISTP. For d, only IMP-8 magnetometer CDFs appear because the magnetometer was chosen. The user is presented with a hierarchy of choices, where each choice limits the following choice to only those CDFs of interest. For this example, a single CDF per panel per plot was chosen in Figure 3. If a multi CDF option was chosen, then more than one choice is allowed on following panels. The final selection, e, is of the key parameters to be plotted on each panel. The user is presented with only the choices that follow from those made before. In this example, 3 panels were chosen with a single CDF per panel per plot so it possible to choose only from this one CDF. The key parameters available are all listed so



that users do not have to know what variables are in the file, or the names of the variables. It is possible to customize the plot from the graphics setting menu, as shown in Figure 3f. The user selects the plot style, the axis scaling (linear or log), and the number of tick marks. It is not necessary to look at *f* unless the data passes a day boundary; defaults are chosen if no other selection is made.

Section C, shown in Figure 3, supplies information such as the Standard Formatted Data Unit (SFDU) label or the skeleton table associated with a particular CDF. The SFDU label contains the global metadata from a CDF, and the skeleton table (discussed in section 2.1) lists all variables and their associated attributes as well as the global metadata. Users can obtain statistics on a particular CDF data set with a call to a CDF toolkit routine. It is also possible to print or erase the screen containing the graphics. Section D in Figure 3, controls the actions such as to plot, to rescale, etc. Figure 3g shows options for rescaling the plots.

Underneath the user view is code that takes advantage of the global and variable attributes which are standard with the ISTP Guidelines. Users are isolated from the format of the data and so can plot data without any knowledge of CDF or the ISTP Guidelines. Defaults are chosen from the attributes associated with the variables in the CDF: LABLAXIS, UNITS, SCALEMIN, and SCALEMAX. If LABLAXIS and UNITS are not included in the CDF, "no such entry" will appear in their place. For this example, the resulting plot using IMP-8 magnetometer data is shown in Figure 4. This data is from 7 July 1993 and was produced with version 1.6 of the Key Parameter Generation Software (KPGS). The magnetic field measurements are shown in GSM polar coordinates. The labels and units, as well as the scales, came from the variable metadata provided in the CDF. The magnetic field data was provided as a vector (a single variable with 3 components) but the components and their associated metadata were split apart and plotted separately. Later releases of the software will include use of other attributes: FILLVAL, DEPEND_0, SCALETYP, AVG_TYPE, OFFSET_0, DELTA_PLUS_VAR, and DELTA_MINUS_VAR.

ISTP Key Parameter Visualization Tool

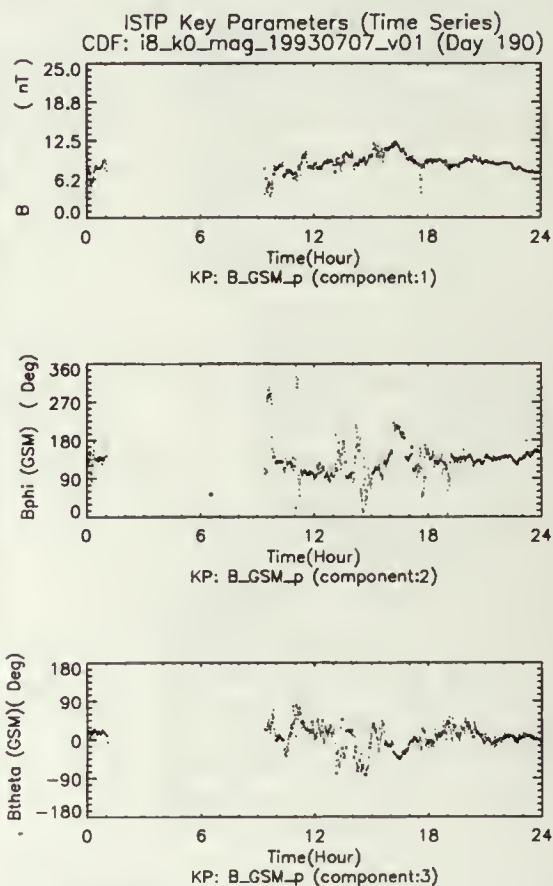


Figure 4. IMP-8 magnetic field key parameter data from 7 July 1993.

4. SUMMARY AND CONCLUSIONS

The International Solar-Terrestrial Physics Program will combine existing and future spacecraft and ground-based measurements from four agencies: NASA, ESA, ISAS, and IKI spread over many countries. This approach to global science places special emphasis on data standards in order to provide easy and effective data sharing among the Space Physics community. The Common Data Format (CDF) with the ISTP Guidelines provides transportable data sets and software and isolates the end user from low level programming. The ISTP Key Parameter Visualization Tool allows the user to plot data conforming to the ISTP CDF standard

without needing outside information or detailed descriptions of the data.

Acknowledgments. The authors would like to thank the many people who contributed to this work including R. E. McGuire, W. H. Mish, G. Goucher, and R. Burley. We especially thank R. Lepping for providing the IMP-8 magnetometer data shown in this paper. We would also like to thank R. Kilgore and J. S. Friedlander for artwork and photography. We acknowledge the Space Science Data Operations Office (SSDOO), and within that the Space Physics Data Facility (SPDF) and the National Space Science Data Center (NSSDC), and the International Solar-Terrestrial Physics Program (ISTP) for support for this project.

REFERENCES

- NSSDC CDF User's Guide for VMS or UNIX Systems, NSSDC/WDC-A-R&S 91-31, Version 2.3, 1992. (A copy may be requested via SPAN — NSSDCA::CDFSUPPORT or INTERNET — CDFSUPPORT@NSSDCA.GSFC.NASA.GOV.)
- R.L. Kessel, R.E. McGuire, H.K. Hills, Jr., J. Love, ISTP Guidelines for CDF, Appendix A in ISTP KPGS Standards & Conventions, Feb., 1994. (Request copy through the ISTP Project at NASA Goddard Space Flight Center.)

Visualization Software for a Global Three-Dimensional Upper Atmosphere Model

**BENJAMIN T. FOSTER, RAYMOND G. ROBLE,
E. CICELY RIDLEY**

*High Altitude Observatory, National Center for
Atmospheric Research, P. O. Box 3000, Boulder, CO*

A software environment has been developed to assist in analysis and visualization of numerical outputs from NCAR's thermospheric general circulation models. The codes produce vector graphics, raster imagery, and time-dependent animation of model results, with provision for comparison with empirical models and relevant data from satellite and ground based instruments. Separate programs are executed on the NCAR CRAY YMP8/864 and HAO divisional Sun workstations, providing both large volume high speed batch and custom interactive computing environments. The visualization system assembles a suite of existing software packages including NCAR Graphics, IDL, netCDF, and the Xt intrinsics library with the Athena widget set.

1. MODEL INTRODUCTION

The principal object for which the visualization system has been developed is the thermospheric general circulation model (TGCM) and its several descendants. The original National Center for Atmospheric Research (NCAR) TGCM was introduced by Dickinson et al. (1981) as a three-dimensional, time-dependent numerical general circulation model of the thermosphere. Over the past decade, the model has been extended to include the effects of upward propagating tides from the middle atmosphere, auroral particle precipitation, and self-consistent mutual coupling between the thermospheric neutral gas and ionospheric plasma (thermosphere-ionosphere general circulation model, or TIGCM). In 1991, an interactive dynamo model was introduced to calculate self-consistent electrodynamic interactions between the thermosphere and ionosphere [Richmond et al., 1992], allowing calculation of global electric potential distribution and ion drift velocities (thermosphere-ionosphere-electrodynamics general circulation model, or TIE-GCM). Most recently, a version of the model has been extended downward into the mesosphere and upper stratosphere (thermosphere-ionosphere-mesosphere-electrodynamics general circulation model, or TIME-GCM) [Roble and Ridley, 1993].

These models use an effective 5 degree global latitude/longitude grid with 25 or 45 constant pressure surfaces in the vertical, for the TIE-GCM and TIME-GCM, respectively. The model time step is typically 5 minutes, with model histories being saved typically every hour for a full day simulation. The model is run on the NCAR CRAY Y-MP8/864 (8 processors, 64 MW central memory with SSD), using about 23 minutes of CPU time for a one-day TIE-GCM simulation. Multiple histories are stored

in history volume files on the NCAR Mass Store System (MSS) (the principal storage medium is IBM 3480 tape cartridges of 200 Mb each). Full sized history volume files are about 170 Mb. Fifteen fields are saved on TIE-GCM histories, and 30 fields on TIME-GCM histories.

2. POST-MODEL PROCESSING

Post-model processors developed for visualization obtain CRAY binary model history files from the NCAR MSS. Codes executing on the CRAY read these histories directly and call NCAR Graphics to produce contours or surface plots of two-dimensional slices selected from the three-dimensional model grid. Because the large raw history files are awkward to download and read with IEEE Sun machines, the CRAY codes are also capable of writing selected portions of the histories as netCDF files, which may subsequently be used by the Sun processors. The Sun codes operate in an X-Window environment and use IDL to produce two-dimensional raster images (Figures 5, 6, and 7). NCAR Graphics is used to produce monochrome or color-fill contour maps (Figures 1-4 and 8-10). The same netCDF history files may be read by both the IDL and NCAR Graphics processors, and may also be used with other software being developed by the visualization community at large.

Capabilities of the post-model processors may be described in five functional categories:

1. snapshot;
2. time-dependent;
3. animation;
4. difference fields; and
5. model-data comparison.

Snapshot codes display model results at a single universal time (UT) (one history). This includes contouring fields over 11 possible mapped projections of the globe (e.g., cylindrical equidistant in Figures 1, 3, and 5; satellite view in Figures 2, 4, and 8; and polar stereographic in Figure 7), zonal and meridional vertical cross-sections (e.g., Figure 6), and vertical profiling of global means. Linear or log interpolation of fields from the pressure coordinate system of the model to constant height surfaces is also a capability of the snapshot processors.

Time-dependent codes display model results from multiple UTs, typically one-to-ten days with

hourly resolution. Fields are contoured with UT and solar local time on the X-axis, and latitude, pressure, or height on the Y-axis (e.g., Figure 10). These plots help visualize the model response to time varying quantities such as those found during geomagnetic storm events. Time-dependent vertical profiles at selected grid point locations ("station processors") are often used for comparison of model results with ground-based instruments.

Time-dependent analysis of the model is assisted by animation of time-interpolated results. For example, a model field may be interpolated between hourly histories, and contoured over a satellite view projection with each frame interval representing 20 minutes of model time and 5 degrees rotation of the earth. If this frame series is recorded to video tape at a rate of 10 frames per second, a ten-day model simulation may be viewed in a 72 second video animation. Color raster image and monochrome or color-fill contour animations have proved useful in analyzing time-dependent phenomena in the models. The NCAR Text and Graphics System (TAGS) is capable of recording image sequences in a variety of video formats including VHS, SVHS, Betacam-SP, and Umatic-SP. Before recording finished sequences to video tape, animations may be previewed on local workstations using the Ximage software developed at the National Center for Supercomputing Applications (NCSA).

Difference field codes obtain results from two separate model runs and display raw or percentage differences of selected fields. This is useful for isolating the effects of altering model input quantities, or evaluating model predictions under different geophysical conditions.

Several independent codes have been developed for comparison of the TGCM models with empirical models and instrument data. Neutral atmosphere simulations are compared with the mass spectrometer and incoherent scatter model (MSIS) [Hedin, 1991], and ionospheric parameters are compared with the International Reference Ionosphere (IRI) [Belitiza, 1986, 1990]. Several studies have been made with ground based incoherent scatter radars, as well as with satellite programs including the Air Force SETA satellites, the Dynamic Explorer Satellites (DE-I and DE-II), and the Upper Atmosphere Research Satellite (UARS).

3. CONCURRENT-MODEL PROCESSING

Often model development and analysis requires knowledge of internal parameters calculated during a model run, but not preserved in history files following model completion. To allow display and evaluation of these parameters the model itself is introduced into the visualization software and allowed to start up from a previously written history and execute a single time step. These programs execute the model code as a subprogram and are termed concurrent model processors. After the model has completed a time step, the fields are displayed in the same manner as the post-model processors. An advantage of this approach is that code transfer from the model to the processors is unnecessary, eliminating the potential for diverging versions of model and processor to produce conflicting results. Concurrent model processors assist in diagnostic analysis of parameters that exist only during model execution such as heating and cooling rates, conductivities, ion drag forcing, and several terms from the momentum and thermodynamic equations used in the models.

Acknowledgments. The National Center for Atmospheric Research is sponsored by the National Science Foundation. IDL (Interactive Data Language) is a product of Research Systems, Inc. (Boulder, CO). NetCDF (Network Common Data Form) is a product of Unidata (Boulder, CO).

REFERENCES

- Bilitza, D., International reference ionosphere: Recent developments, *Radio Sci.*, 21, 343-346, 1986.
- Bilitza, D., International reference ionosphere 1990, National Space Science Data Center/World Data Center A for Rockets and Satellites, *NSSDC/WDC-A-R&S 90-22*, Nov. 1990.
- Dickinson, R.E., E.C. Ridley, and R.G. Roble, A three-dimensional general circulation model of the thermosphere, *J. Geophys. Res.*, 86, 1499-1512, 1981.
- Hedin, A.E., Extension of the MSIS thermosphere model into the middle and lower atmosphere, *J. Geophys. Res.*, 96, 1159-1172, 1991.
- Hedin, A.E. et al., Revised global model of thermosphere winds using satellite and ground-based observations, *J. Geophys. Res.*, 96, 7657-7688, 1991.
- Richmond, A.D., E.C. Ridley, and R.G. Roble, A thermosphere/ionosphere general circulation model with coupled electrodynamics, *Geophys. Res. Lett.*, 19, 601-604, 1992.
- Roble, R.G., E.C. Ridley, A.D. Richmond, and R.E. Dickinson, A coupled thermosphere/ionosphere general circulation model, *Geophys. Res. Lett.*, 15, 1325-1328, 1988.
- Roble, R.G., and E.C. Ridley, A thermosphere-ionosphere-mesosphere-electrodynamics general circulation model (TIME-GCM): Equinox solar minimum circulation (30-500 km), *Geophys. Res. Lett.*, 19, submitted, 1993.




Figure 1. Color fill contours of total density ($O+O_2+N_2$) at 300 km altitude mapped onto a cylindrical equidistant global projection. The graph at bottom shows the geophysical index K_p during a ten day geomagnetic storm period in March, 1979. The time of the density image is at the end of the red portion of the K_p graph. This figure is a still frame from a video animation of a ten day model simulation.




Figure 2. Color fill contours of atomic oxygen at approximately 120 km, mapped onto a satellite view projection. K_p graph as in Figure 1. This plot shows depletion of atomic oxygen at high northern latitudes during the March, 1979, geomagnetic storm period. This is also a still frame from a video animation.

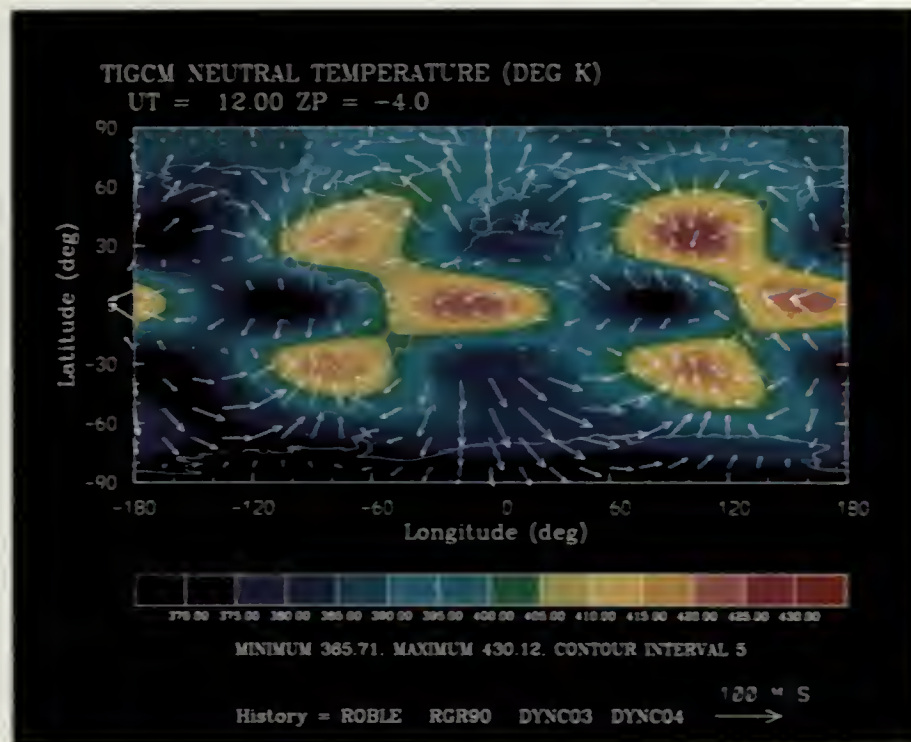


Figure 3. Color fill contours of neutral temperature on a cylindrical equidistant projection, with overlying vectors of neutral wind velocities. The effect of semi-diurnal upward propagating solar tides are visible in the temperature contours at this altitude (approximately 120 km).




Figure 4. Neutral temperature and wind response to the July, 1991, total solar eclipse at an altitude of approximately 300 km. These are difference fields between two TIE-GCM runs: one with the eclipse subtracted by one without the eclipse. The solid line shows the path of totality. Vectors indicate winds converging into the area cooled by the eclipse's shadow.

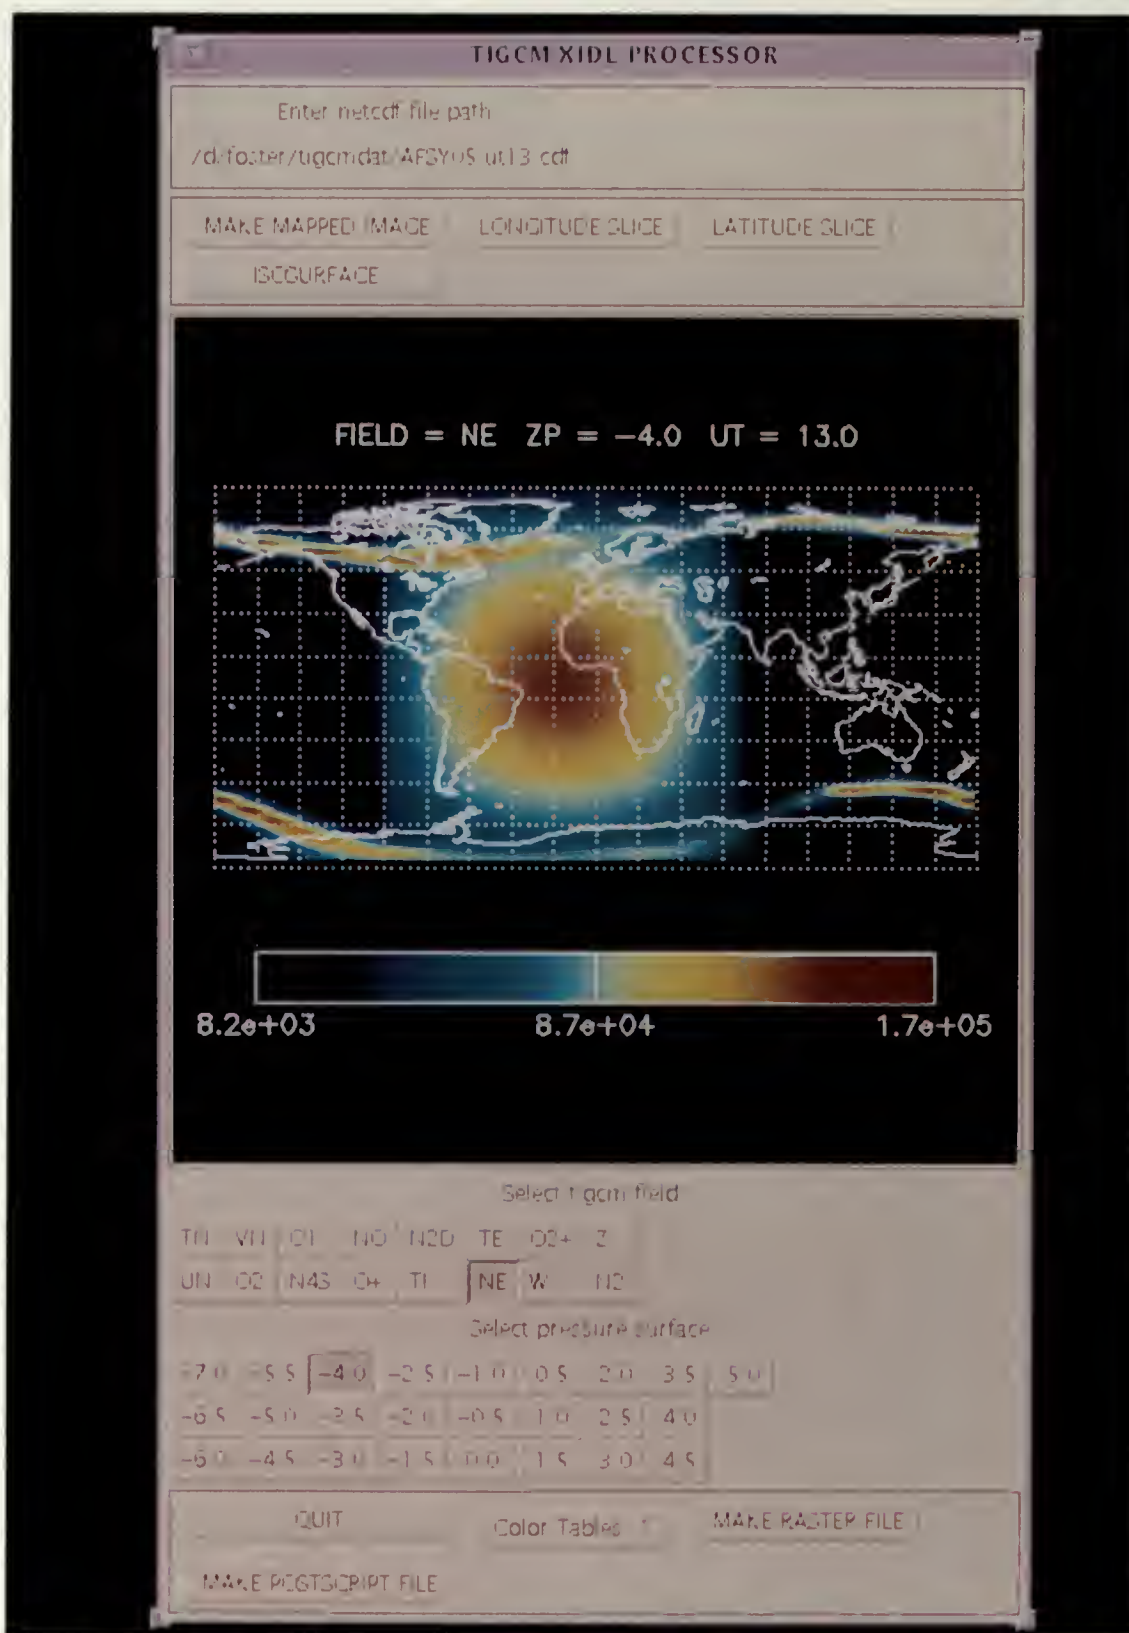


Figure 5. Main control window of an interactive snapshot Sun processor running IDL with OpenWindows. The draw widget shows a raster image of electron density at the -4 pressure surface (about 120 km) on a cylindrical equidistant projection. Density highs are seen at solar local noon (near center of the image), and at high latitudes (the auroras).




Figure 6. Raster image with overlaid contours showing a vertical slice of electron density along the 57.5 degree north latitude circle. This image was produced by the snapshot Sun processor running IDL (as in Figure 5). Local noon is at zero degrees longitude. The vertical axis is in a log pressure scale. Pressure level -7 corresponds to 95 km, and pressure level +5 (top of the model) is typically 400 km during solar minimum conditions and 600 km during solar maximum conditions.




Figure 7. Favorable model-data comparison of the typical dual vortex wind pattern at high latitudes near 300 km altitude. This is a stereographic projection with a raster image of neutral temperature. Predicted neutral wind velocities (white arrows) are from the TIE-GCM, and observed wind velocities (black arrows) were measured during an orbit pass of the Dynamic Explorer-II satellite.




Figure 8. Color fill contours of electron density (pressure level -4, approximately 120 km) on a satellite view projection showing the effect of the aurora in the northern hemisphere. This plot was produced by the snapshot post-model processor calling NCAR Graphics.




Figure 9. Stack of five color contour maps showing nitric oxide densities at different vertical pressure levels (Z_p -6 to -2 corresponding to approximately 100 to 150 km altitude). Each plane is a global latitude versus longitude map. This plot shows the vertical variability of high latitude nitric oxide density. Numbers to the right under the Z_p values show the range of nitric oxide (\log_{10}) at each level.




Figure 10. Atomic oxygen density difference fields (at approximately 300 km) over an 8 day period along the 70 degree west longitude (latitude on the Y-axis). A strong geomagnetic storm event commenced during the third day. Decreased atomic oxygen densities are predicted at high latitudes throughout the storm (see also Figure 2).

A Prototype Upper Atmospheric Research Collaboratory (UARC)

C.R. CLAUSER¹, D.E. ATKINS², T.E. WEYMOUTH³,
G.M. OLSON⁴, R. NICIEJEWSKI¹, T.A. FINHOLT⁴,
A. PRAKASH⁵, C.E. RASMUSSEN¹, T. KILLEEN¹,
T.J. ROSENBERG⁵, D. DETRICK⁵, J.D. KELLY⁶,
Y. ZAMBRE⁶, C. HEINSELMAN⁶, P. STAUNING⁷,
E. FRIIS-CHRISTENSEN⁷, S.B. MENDE⁸

The National Collaboratory concept has great potential for enabling "critical mass" working groups and highly interdisciplinary research projects. We report here on a new program to build a prototype collaboratory using the Sondrestrom Upper Atmospheric Research Facility in Kangerlussuaq, Greenland and a group of associated scientists. The Upper Atmospheric Research Collaboratory (UARC) is a joint venture of researchers in upper atmospheric and space science, computer science, and behavioral science to develop a testbed for collaborative remote research. We define the "collaboratory" as an advanced information technology environment which enables teams to work together over distance and time on a wide variety of intellectual tasks. It provides: (1) human-to-human communications using shared computer tools and work spaces; (2) group access and use of a network of information, data, and knowledge sources; and (3) remote access and control of instruments for data acquisition. The UARC testbed is being implemented to support a distributed community of space scientists so that they have network access to the remote instrument facility in Kangerlussuaq and are able to interact among geographically distributed locations. The goal is to enable them to use the UARC rather than physical travel to Greenland to conduct team research campaigns. Even on short notice through the collaboratory from their home institutions, participants will be able to meet together to operate a battery of remote interactive observations and to acquire, process, and interpret the data.

¹University of Michigan, Space Physics Research Laboratory, Ann Arbor, MI

²University of Michigan, School of Information and Library Studies, 304 W. Engineering, Ann Arbor, MI

³University of Michigan, Department of Electrical Engineering and Computer Science, Ann Arbor, MI

⁴University of Michigan, Cognitive Science and Machine Intelligence Laboratory, 701 Tappan, Ann Arbor, MI

⁵University of Maryland, Institute for Physical Science and Technology, College Park, MD

⁶SRI International, Geoscience and Engineering Center, 333 Ravenswood Ave., Menlo Park, CA

⁷Danish Meteorological Institute, Solar Terrestrial Physics Division, Copenhagen, DK

⁸Lockheed Palo Alto Research Laboratory, 3251 Hanover St., Palo Alto, CA

1. INTRODUCTION

We describe here a multidisciplinary effort linking research in computer science, behavioral science, and upper atmospheric and space physics. The purpose of this effort is to conceive, develop, deploy, test, evaluate, and integrate a high performance group centered computer environment into collaborative experimental activities ongoing in the space science research community. Such group computing environments, we predict, will become an important component of the National Information Infrastructure (NII) initiative, which is envisioned as the high performance communications infrastructure to support future national scientific research.

Because upper atmospheric and space science is influenced by a broad and diverse set of regions and processes, progress requires an ever increasing synthesis of information from a wide variety of experimental data and the interaction between scientists in varying areas of research. Thus, the computing infrastructure to support this research is becoming increasingly necessary to support collaborative efforts to acquire and synthesize information. The technology to enable such interactions is now developing rapidly and the United States is making a major national commitment to develop and deploy such technology.

By way of background, the term telescience was coined in 1985 to capture the notion of an investigator participating in experimental operations at a distance supported by an electronic infrastructure. Such a notion was developed to characterize some types of experimental operations envisioned for the space station. A variety of scenarios were envisioned in which the principal investigator on an experiment was not able to be

physically present on the space station, but whose critical knowledge and insight were necessary to the success of the experiment. Thus, a virtual presence for the investigator, working with the space station crew, is the logical requirement for the success of these experiments.

The Sondrestrom Upper Atmospheric Research Facility in Kangerlussuaq, Greenland is a remote ground facility, jointly supported by the National Science Foundation and the Danish Meteorological Institute and maintained and operated by SRI International [Kelly, 1983; Wickwar et al., 1984; Clauer et al., 1984]. This facility was proposed to NASA in 1989 to be a good analogue to develop, test, and evaluate the tools necessary for telescience. The Sondrestrom facility could be utilized to test a variety of remote experimental activities which may be similar to those envisioned for a space station. For example, the facility contains complex equipment which is maintained and operated by a site crew, is located in a remote site with limited logistical support and limited travel access, and often requires that the investigators travel to the site to undertake experimental operations. Remote access to support experimental operations could have many benefits to the science undertaken at the Sondrestrom facility, as well as provide an effective testbed in which to develop the tools to enable such interactions.

A telescience testbed program centered upon the Sondrestrom incoherent scatter radar was supported by NASA for two years. With this support, an electronic satellite link to the facility was established and the development of X-Windows based software to provide remote display of real time data from the radar was developed. The Sondrestrom facility, however, contains a variety of separate, but often synergistic, experiments operated by a variety of scientists who are distributed around the world. Thus, the telescience infrastructure developed through the NASA supported telescience testbed has provided the enabling capability for a much more ambitious activity: the Upper Atmospheric Research Collaboratory (UARC) testbed.

The notion of telescience was greatly expanded in 1989 with the concept of the National Collaboratory [National Collaboratories, National Academy Press, Washington, DC, 1993]. The

National Collaboratories report, prepared under the auspices of the National Research Council, concluded that "collaboratory testbed programs have the potential to address important scientific needs while simultaneously representing a key step toward developing national and global infrastructures." The collaboratory is described as a "center without walls in which the nation's researchers can perform research without regard to geographical location — interacting with colleagues, accessing instrumentation, sharing data and computational resources, and accessing information from digital libraries." Thus, the collaboratory infrastructure is envisioned to support distributed interaction between people, access to remote information sources and digital libraries, and access to and interaction with remote and unique facilities.

It is the aspects of access and interaction with remote and unique facilities, as well as distributed interactions between people that are the focus of the prototype testbed which we have named the Upper Atmospheric Research Collaboratory (UARC). Beginning in September, 1992, the Computer Information Science and Engineering Directorate (CISE) in cooperation with the Atmospheric Sciences Directorate of the National Science Foundation has funded this major collaboratory testbed within the space science community. It was felt that progress would be accomplished most rapidly by using a testbed approach formed around a collaborating group in a real world environment. Thus, the UARC has been formed around the ongoing research activities among a group of space scientists engaged in experimental operations at the Sondrestrom Upper Atmospheric Research Facility in Greenland. The UARC is being created using a user-oriented, rapid prototyping research approach at the University of Michigan, SRI International, and the other testbed sites — the Danish Meteorological Institute, the University of Maryland, and the Lockheed Palo Alto Research Laboratory. In this project, an interdisciplinary group of investigators is supported to conduct coordinated experimental research related to creating and evaluating distributed environments to support team science. The goal is to build a networked environment to support interactive observational campaigns and team science using multiple instruments and a distributed group of investigators located at their

home institutions rather than at the Sondrestrom facility.

2. THE SONDRESTROM COLLABORATORY TESTBED

The instrumentation presently supported to participate in this prototype environment, their principal investigators, and their institutions include: (1) incoherent scatter radar, John Kelly, SRI International; (2) imaging riometer, Peter Stauning, Danish Meteorological Institute and Ted Rosenberg, University of Maryland; (3) magnetometers, Eigil Friis-Christensen, Danish Meteorological Institute; (4) Fabry Perot interferometer and optical spectrometers and photometers, Rick Niciejewski and Tim Killeen, University of Michigan; and (5) all-sky imaging television camera, Steve Mende, Lockheed Palo Alto Research Laboratory.

We plan for the testbed to evolve through three design phases: The “wire service” phase, the “point and talk” phase, and the “fully shared control” phase. We are presently in the wire service phase, where, using existing technology, the instruments at Sondrestrom generate a stream of data that is collected, stored, forwarded, and

displayed at user sites. The point and talk phase is an extension to the wire service phase that allows researchers to interact jointly while they are located at several distributed sites. We have entered this phase with the implementation of a set of simple communications windows that broadcasts discussion to all interactive users. The fully shared control phase builds upon the point and talk phase to provide more natural and fluid group interactions using emerging voice and video communication technologies. The fully shared control phase also supports the remote control of the instruments.

Figure 1 provides a schematic view of the testbed showing both initial configuration at Sondrestrom (a), and the user’s laboratories (b), and also showing the technology which we hope to achieve and, in part, operationalize during the testbed (c). In the first level of communications represented by (b), a user may log onto a computer in Sondrestrom or a local server located elsewhere which is itself logged onto a computer in Sondrestrom via the Internet. Data may be transferred via standard file transfer techniques for display. This is being automated to become the wire service phase now. The project will evolve to the level of communication illustrated in (c). Here we show multiple users showing multiple data in

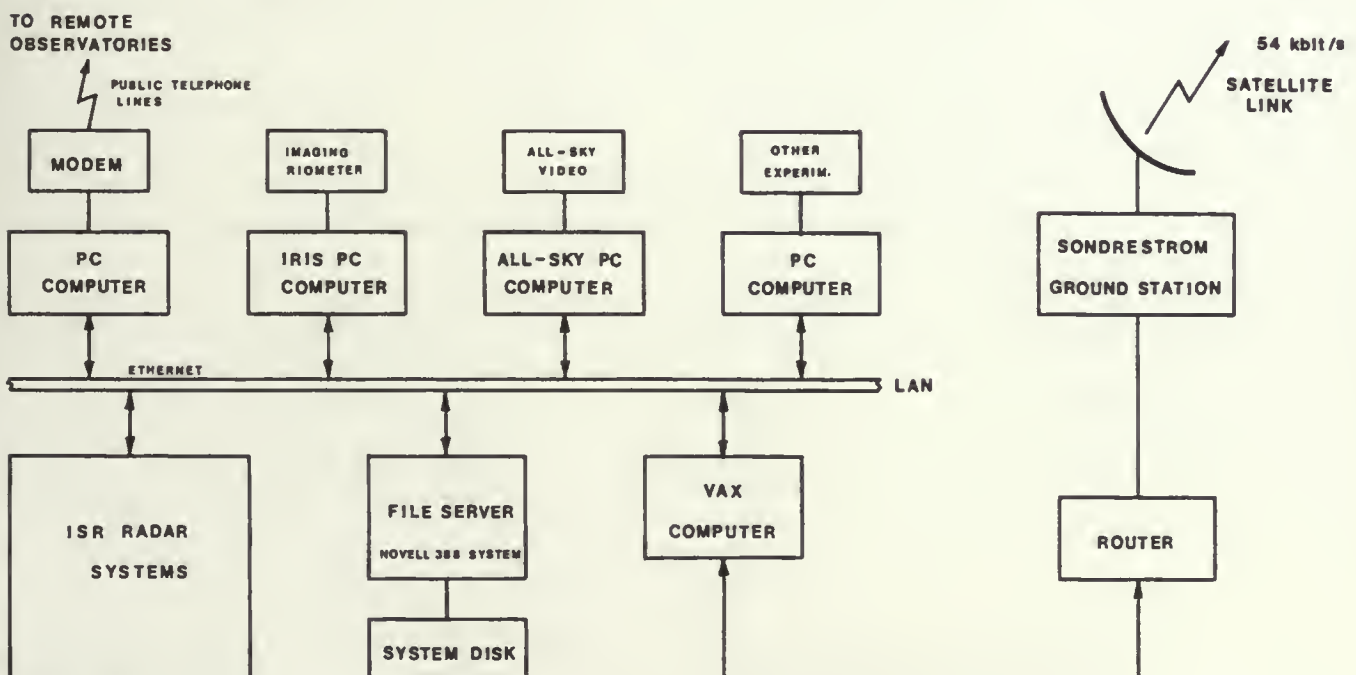


Figure 1a.

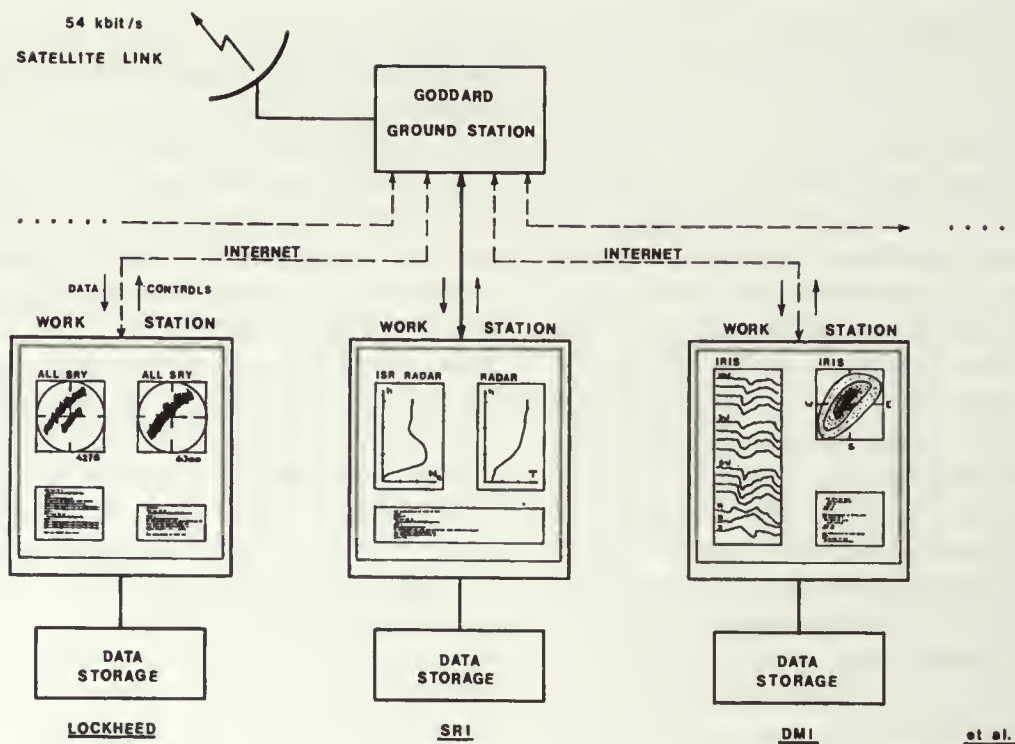


Figure 1b.

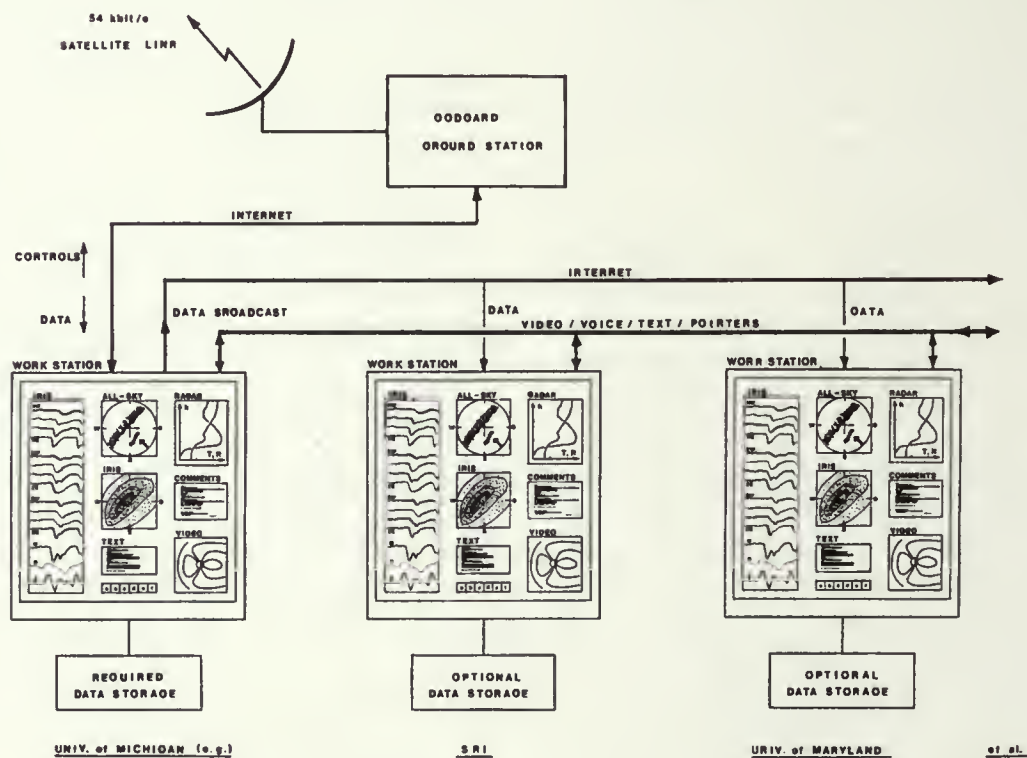


Figure 1c.

windows which support interactive pointing, annotation, sketching, text, voice, etc. Control panels will also permit investigators to interact with their instruments while also providing safeguards to prevent conflicting or inappropriate actions. Parts of (c) have been implemented, including the ability for multiple users to view data in multiple windows and communicate through a shared text window.

Because the project involves a relatively small number of institutions and users, we have chosen the luxury of utilizing a homogeneous computing environment and to not deal with interoperability issues. This permits us to focus on the issues directly related to design specifications for a collaborative environment. We have chosen to use the NeXT computing environment including the NeXTStep Interface Builder. NeXTStep was chosen because it tightly integrates both the operating system, supporting software, and development environment under a unifying object-oriented software library. NeXTStep also allows instances of software objects running on one workstation to be seamlessly distributed to other workstations on the Internet. This supports the modular development of distributed group-oriented software in an evolutionary fashion. It represents, in our view, the best existing example of the type of computer system which will be common in the future. Also very important to our approach is that the use of this homogeneous, object-oriented environment has enabled us to pursue a rapid prototyping approach to system design and testing. We envision a cycle time of about 3 months for software design, development, deployment, evaluation, and revision. Development can proceed much faster using the NeXT development system than systems, such as X-Windows, for example. It is our opinion that greater efficiency will be achieved by doing the development in the NeXT environment and then porting the developed software to X if this becomes desirable at a later time. Also, since distributed objects are now supported separately in the UNIX environment, it will be possible to continue to utilize the NeXT server model with distributed objects passed to X-Window clients for display. Such a strategy will permit the further extension to a wider community in an efficient way since only the display software will have to be rewritten for the X-Window clients.

In Figure 2 we show a screen display from the current software captured from a past experimental operation. The various windows on the display show the data acquired by the radar—line of sight velocity (bottom right) and electron density (top right) from an azimuth scan. Also shown are the communications windows—one for entering the user's messages (top left), and the other one which displays the messages sent by all users (bottom left). While there are threads of different conversations which develop in the message window, this has not been a problem and users have been able to ignore or participate in conversations according to their desired involvement. Another data display window showing ionospheric plasma density as a function of time and altitude is partially visible lying underneath the two azimuth scan data display windows. The control menu window for the program is shown in the upper left corner and icons for other NeXT tools and applications are shown along the right side of the display.

In Figure 3 we show the extent of the present UARC network. In all, 14 workstations are supported and are distributed between developers and social scientists at the University of Michigan and the space science users who themselves are distributed between Denmark, Greenland, Maryland, Michigan, and California. The various instruments in Sondrestrom are controlled by PCs and are connected to the Sondrestrom local area network (LAN). A VAX computer and a router provide connection to the external TCP/IP internet.

As prototypes of various capabilities emerge, they will be tested and evaluated. The tools will develop through a series of iterations whereby feedback from the evaluations define new requirements for the tools. As prototype tools are refined and their utility demonstrated, they will become operational and will be maintained on the foundation layer of the testbed.

A fundamental characteristic of our approach is a rapid prototyping cycle involving the definition of user requirements, prototyping, deployment, testing and evaluation, and revision. It is our goal to operate this cycle typically on a 3 or 4 month cycle time. Using the NeXT and an object-oriented approach to design and implementation, we have been able to "turn over" this cycle twice in the past six months. Specifically, the initial version of the



Figure 2.

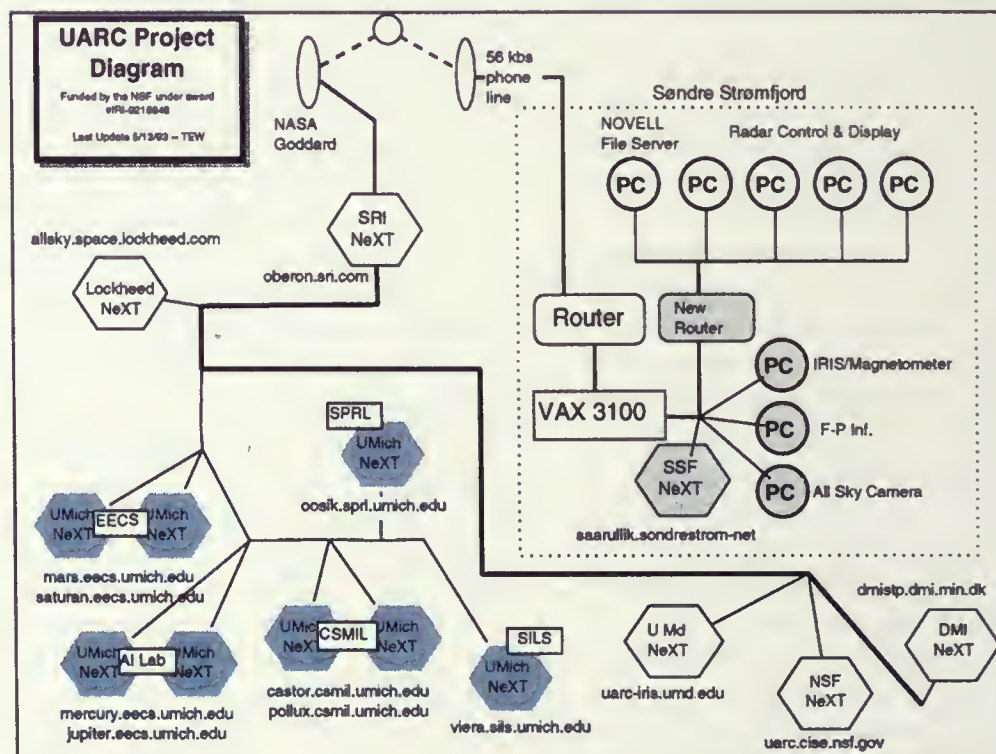


Figure 3.

UARC Project software provides radar data distribution (from Greenland to anywhere in the Internet) and 2 types of data displays. That version was developed in 5 weeks. The performance of this first version was evaluated during a radar campaign in April 1993. Utilization of the software was surprisingly robust. It supported, at times, over 12 sites including ones in Denmark, Greenland, and in the US, all observing the operations of the radar and interacting over choices of the radar operation modes. Essentially, the same version of the software has been demonstrated publicly at the American Geophysical Union Spring Meeting in May [Clauer et al., 1993]. User reactions to that first version were then incorporated into a major revision of the software, which was used to support a series of radar experiments in June, 1993.

An important and unique aspect of this multidisciplinary effort is the behavioral science research. Behavioral scientists are involved in this project in two ways. First, they are strongly involved in the software design. The behavioral scientists have taken the lead to define the objects which the space science users utilize and assist the computer scientists to codify these objects in our object oriented programming approach. The behavioral scientists are directing the object oriented design methodologies and are monitoring the use of the software to assist in the iterative redesign of the system. It is important for this project to develop generic design specifications which may have wide spread application to electronic collaboration across disciplines. The behavioral science team has a strong history in the development and evaluation of object oriented software to support team engineering efforts. Second, the behavioral scientists are providing documentation regarding the effect of introducing this new technology into the scientific practice of the space scientists who are using the testbed. Concurrent with the startup phase, the behavioral scientists have collected information about current work practices among the space scientists at the various sites. These measurements will provide a basis upon which to assess the effects of the new tools as they are implemented, tested, and evaluated in the testbed. An important outcome of the project will be a quantitative documentation of the benefit or the changes which result from the utilization of the collaboratory.

There are a number of interesting technical issues associated with providing the kinds of capabilities needed by the Sondrestrom users. To restate what we said earlier, the three classes of capabilities we are providing are: sharing of data from the Greenland instruments in real time over a wide area network; control of the instruments over the network; and collaboration tools that allow the scientists to work together over the network. We are developing a framework within which we can explore these technical issues and evolve tools that will work under the operating conditions found in the space science research (e.g., complex data displays and widely separated workstations). Our approach is a mixture of an object-oriented representation of the basic constructs (e.g., data, instruments, analysis procedures, data visualization tools) and a data flow model for handling the coupling between instruments and user interactions.

The need for an explicit model of the user interaction becomes clear when we consider, for example, one of the key collaboration capabilities we plan to offer—the ability to share annotations on the data from the instruments. How should the annotations be represented? Since each user can choose to display data from a particular instrument in different ways, it seems like the best representation is to associate the annotations with the data itself rather than with the visual display. This way each user can examine the annotations in the context of the particular way they have chosen to display the data. Of course, in some cases the annotations may not make sense unless the display is of a particular type. But, then someone reading the annotation could easily switch to that type of display, and the annotation would still be associated with the relevant data item.

We are exploring how the annotations themselves should be represented and shared. We envision the annotations being of multiple data types themselves, such as text, drawings, or voice. We will represent the annotations in such a way that any degree of sharing is possible, such as read-only or read-write access to others' annotations. These capabilities must be offered so that they work over the wide area network with the typical delays that make precise synchronization impossible.

Other capabilities will be added based upon the emerging needs of the user community, as well as

the enabling technologies produced by the computer science research team. The goal of the basic computer science and engineering (CSE) research in the project is to explore new paradigms (with appropriate metaphors) for tools supporting group interaction over real-time data. The CSE group is exploring the structure of toolkits, window systems, and underlying support for sharing, migration, and selective replay of windows and collaborative sessions.

3. DISCUSSION

The collaborative technology developed through the research undertaken within the UARC project should be widely extendable through the space science community, as well as to other areas of science and, perhaps, beyond. Since much of the collaboration technology is expected to also have generic value, the development here may also provide future benefit to the design and use of similar collaboratories in other areas, such as education, engineering, and business.

As the project progresses, we intend to make it more comprehensive by adding additional instruments and users. As we consider the possibilities for extension of the UARC, issues of interoperability become more important. Since the NeXTStep operating system is now available for Intel 486 processors, it is relatively easy and inexpensive to join the homogenous UARC development environment. We note also that Hewlett Packard and SUN have announced they will support NeXTStep on their high performance RISC workstations. Still, we realize this is not a suitable solution for general expansion.

Since distributed objects are now supported independently in the UNIX environment, it is possible to utilize the UARC NeXT server in conjunction with an X-Windows display environment at individual user sites. The only development which is required for this to happen is the data display software for the X-Windows environment. The UARC distributed objects could be passed to the X-Windows system, or any other system

supporting distributed objects, and the display software for the objects can be developed separately. This allows us to maintain the foundation of the collaboratory and the rapid prototyping development environment with a clean separation from the user display environment. It is this approach that we feel will be the most effective and efficient for more general future extension of the UARC to additional users.

Acknowledgments. Support for this work has been provided by the National Science Foundation through a cooperative agreement IRI-9216848 to the University of Michigan and research has been supported by grant ATM-9106958. The Sondrestrom Upper Atmospheric Research Facility is operated by SRI International for the National Science Foundation through cooperative agreement ATM-8822560.

REFERENCES

- Clauer, C.R., D.E. Atkins, T.E. Weymouth, G.M. Olson, R. Niciejewski, T. Finholt, A. Prakash, C.E. Rasmussen, T.J. Rosenberg, J.D. Kelly, Y. Zambre, P. Stauning, E. Friis-Christensen, and S.B. Mende, A prototype upper atmospheric research collaboratory (UARC), *EOS, Trans. Amer. Geophys. Union*, 74, 267, 1993 (abstract).
- Clauer, C.R., P.M. Banks, A.Q. Smith, T.S. Jorgensen, E. Friis-Christensen, S. Vennerstrøm, V.B. Wickwar, J.D. Kelly, and J. Doupnik, Observation of interplanetary magnetic field and of ionospheric plasma convection in the vicinity of the dayside polar cleft, *Geophys. Res. Lett.*, 11, 891, 1984.
- Kelly, J.D., Sondrestrom Radar — initial results, *Geophys. Res. Lett.*, 10, 1112, 1983.
- National Collaboratories*, National Academy Press, Washington, DC, 1993.
- Wickwar, V.B., J.D. Kelly, O. de la Beaujardiere, C. Leger, F. Steenstrup, and C.H. Dawson, Sondrestrom Overview, *Geophys. Res. Lett.*, 11, 883, 1984.

Visualization Tools for the Processing of Airglow Data from RAIDS

GORDON J. MILLER¹, STEFAN THONNARD,
MICHAEL PICONE, CAROLYN MEESUK²

*E.O. Hulburt Center for Space Research, Naval
Research Laboratory, Washington, DC*

In anticipation of large data sets associated with a number of atmospheric imaging instruments being prepared for long term global coverage, NRL is developing graphical interfaces for all aspects of the program. For the first of these projects, RAIDS (the Remote Atmospheric and Ionospheric Detection System), a graphical approach to data handling, visualization, and analysis is envisioned and will set the stage for the satellites that follow. An overall system of hardware and a set of software "tools," that will allow for both the routine handling of all data and the analysis of large data sets assembled by scientists and instrument engineers, are currently being developed.

The software for standard processing and visualization of instrument data is independent of computer platform and will allow for easy adaptation from one experiment to another. The processing will produce data sets that have similar characteristics, allowing for easy comparison of data obtained under similar circumstances.

The visualization of both the engineering and scientific data is an important part of the system. By creating graphical environments for engineering evaluations and for scientific analysis data sets can be viewed and analyzed rapidly. This rapid analysis of data will contribute towards a greater portion of the RAIDS data being utilized.

1. INTRODUCTION

The Naval Research Laboratory (NRL) is currently preparing a number of atmospheric imaging instruments to be launched aboard a variety of satellites. The first of these projects, the Remote Atmospheric and Ionospheric Detection System (RAIDS), is due to be launched during 1994. Over the course of its lifetime, RAIDS is expected to produce a data archive of 220 gigabytes or approximately 145 megabytes of compressed data per day. For a volume of data such as this, if scientists wish to analyze any significant portion of the data, methods allowing for processing and analyzing large amounts of data quickly are required.

An important part in the development of analysis software is the creation of meaningful methods of visualizing variations in data over longitude, latitude, altitude, and time. The data analysis is facilitated by a Graphical User Interface (GUI), using "tools" created specifically to process the data from a particular instrument. However, users can easily customize the GUI by creating new graphical tools and including them in a common library. In addition, data sets will be analyzed and recorded on video tape or optical disk, allowing rapid searches of large data sets for events of interest and permitting subsequent presentations.

NRL is in the late stages of developing an automated and graphical data processing and analysis system. This system involves all aspects of the experiment, including mission planning, data processing, experiment evaluations, simulation, and data analysis. By automating the majority of the data handling, engineers can quickly diagnose and correct problems, while scientists can concentrate on data analysis.

¹SFA, Inc., Landover, MD

²Software Technology, Inc., Alexandria, VA

In the following sections, aspects of the NRL system will be explained. The next section will describe the flow of data through the system and how this flow is to be maintained. Then, a description of how coordination between the science and engineering teams will be achieved. Following mission coordination, the ongoing evaluation of the instrument and data integrity, and how this will be performed, will be discussed. The remainder of the paper will be devoted to methods of data handling and data visualization that will be applied to the analysis of the collected data.

2. MISSION PLANNING AND COORDINATION

During the experiment lifetime, coordination between science objectives and engineering concerns will always be an issue. Specific scientific objectives, which should be met, are the coordination of upper atmospheric observations with other experiments and providing observations of specific phenomena at specific times. Along with these science objectives, engineers will be concerned with ensuring the safety and longevity of the instrument, calibration, and instrument sensitivity.

By using software which can predict satellite and experiment parameters, many of the engineering and science objectives can be easily met. By entering orbital and trajectory information, other orbital parameters which are of interest can be computed. These parameters can then be viewed graphically by the scientist for mission planning and the parameters can be saved for later use in computing and planning. Using information gener-

ated by the software, engineers and scientists can graphically see where the instrument will be and what it will be viewing, allowing them to make decisions benefiting both sides.

As the experiment goes on, the engineers must continue to monitor the instrument sensitivities and the satellite attitude. This can be accomplished by using the background stars that will be contained in some of the data. Software will be used to generate the spectrum that the instrument will see from these stars. These spectra can then be compared over time to see how the instrument sensitivity changes over time. Attitude can also be monitored by seeing how the stars progress through the field of view of the instrument.

3. DATA FLOW AND PROMOTION LEVELS

The majority of the NRL system is in place to process and archive the incoming data. This system is written in ANSI - C because of the availability of C compilers, the ANSI standard, the extensive knowledge of C among the programming community, and the suitability of C to the processes of bit shifting and variable size arrays. The purpose of this portion of the system is to ingest data from the experiment and convert it into machine readable format, calculate querying and analysis parameters, archive the data, update the data base, and create "products" which will facilitate the analysis of the data.

Figure 1 shows the overall processing system. The data is received at NRL in a compressed form which requires expansion onto standard byte boundaries before being considered level 1 data.




Figure 1. Diagram showing the flow of data through the system. The arrows go from the processing that is performed to the level of data that is produced.

This data is verified to be in a consistent format to ensure proper operation of the later stages of the processing.

The next stage is the calculation of attributes from the information contained in the expanded data stream. There are a total of 64 attributes which consist of information, such as altitudes, attitudes, angles, and coordinates in different systems. The software itself consists of individual subroutines, written in C, and their calling routines that will compute the information. The calling routines will check the data and call the subroutines to perform the computation of the data. The computation of the attributes will be completed before the next stage of the processing system begins.

At this point, processing will stop and all of the information will be archived. Optical disks have been chosen as the archive medium for their storage capacity per disk. After the data is archived, it remains online for thirty days. The archive medium is then duplicated so that risk of losing data is reduced as low as possible. The duplicates will be used after the data set goes offline to support requests for back data. As the archive is created, all of the attributes that were calculated are used to

create a relational data base that can be used to obtain subsets of data that meet particular criteria. Criteria for subsets can involve any of the 64 attributes that are calculated plus additional information that is generated directly by the instrument. The purpose is to allow the user to obtain data directly and quickly. An example of the querying screen that will be used is shown in Figure 2.

After the integration of the day's data into the data base, images of the instrument data and environmental parameters are constructed. Parts of the instrument data are stripped from the level 1 data and converted into scientific units, producing level 2 spectral data, and the environmental parameters are computed from these level 2 images. These images are produced in a standard format which will allow for easy viewing by scientists without an extensive knowledge of the format and requiring little or no pre-provided software. This part of the system scans through the data for information that matches the search criteria and constructs images out of all the information that meets this criteria. The output from this part of the system is two data files. One contains the actual image of the remote sensing data, while the other

pci_rad_sat_q		Unclassified	
RAIDS PROGRAM CATALOG			
SCAN SEARCH - SATELLITE PLATFORM			
Query Name : <input type="text"/>		Data Source : <input type="text" value="all_raids_scans"/>	
Mode No : <input type="text" value="*"/>		Co-ordinate Type : <input type="text" value="Geodetic"/>	
Satellite Platform (deg)		Line Of Sight (deg)	Time (mm-dd-yy hh:mm:ss)
Latitude	<input type="text" value="0 00"/> to <input type="text" value="50 00"/>	RA	<input type="text"/> to <input type="text"/>
Longitude	<input type="text" value="-67.00"/> to <input type="text" value="-30.00"/>	Dec	<input type="text"/> to <input type="text"/>
Spectrometer Settings		Solar Co-ordinates (deg)	
MUV wavelength	<input type="text" value="2000"/> to <input type="text" value="3000"/> (angstroms)	RA	<input type="text"/> to <input type="text"/>
NUV wavelength	<input type="text" value="3000"/> to <input type="text" value="4000"/> (angstroms)	Dec	<input type="text"/> to <input type="text"/>
NIR wavelength	<input type="text" value="7500"/> to <input type="text" value="8000"/> (angstroms)	Angle to LOS	<input type="text"/> to <input type="text"/>
Alpha Angle	<input type="text"/> to <input type="text"/> (deg)	Geophysical Indices	
***** Use <tab> key to move to the next prompt *****			
Menu(F1) Help(F2) Choices(0) Clear(1) Execute(2) Query(3) >			

Figure 2. Sample querying screen from the RAIDS data base.

contains the information from the data stream that will be used to analyze this data. Attributes are not passed with this image. The information required to compute the attributes is passed with the image and the routines for the computing of the attributes are provided to all the users. This is done to reduce the amount of redundant data that is to be stored. This data will then be archived to magnetic tape and be integrated into the data base as auxiliary data available from a query.

Another important component of the data analysis will be the use of video. By using video, new approaches to data analysis can be developed which will allow individuals to view and analyze time variations in data that would otherwise be difficult to see. Video also allows for easy transport of the data for viewing purposes between locations. This will allow presentations to showcase the actual data that was used for the analysis instead of just a single frame or two.

4. ENGINEERING VISUALIZATION

During the experiment, constant monitoring is necessary. This monitoring will focus on temperatures and voltages on the experiment. If any of these values leave their nominal ranges, a diagnosis and remedial decision must be made. Because of the large amount of information that must be viewed for proper evaluation of the instrument status, methods to allow for easy viewing and monitoring of this information are required.

Two distinct methods have been developed to handle the long-term viewing and the monitoring of this information. At the lowest level, this will consist of routines that will check each value from the instrument against a nominal range, and if any values occur outside this nominal range, it will be reported. By using this tool, engineers will be alerted to situations where the experiment is behaving abnormally. By screening all incoming data, constant supervision is unnecessary.



Figure 3. Sample screen of Engineering Station. Shown are the menus for selecting various components to view and the two windows in which they can be viewed.

Accompanying the continual screening is a graphical engineering station. Once a problem has been identified, engineers must be able to examine the health of the experiment around the time of the event. By using this station, engineers can view either the history of a particular instrument or single values from different instruments simultaneously. With this capability, engineers will be able to quickly diagnose specifically where the problem lies and formulate a solution.

The engineering station, shown in Figure 3, was designed using the IDL X-Windows widget interface library. The station is operated almost entirely by pull down menus, with keyboard interaction being limited to entering frame numbers for viewing. The minimization of keyboard interaction reduces the amount of time that the user must spend in the initialization of the station to view data, and also reduces the probability of crashing the system.

By selecting from various, self-explanatory options under various menus, the user sets up the particular viewing conditions he wants. At this point, he is presented with plot windows for the display of the particular data and a set of scrolling lists to allow for the selection of the information he wants to view. Upon selecting an item from one of the scrolling lists, the information he requested to view is displayed on the drawing area. By looking at the various items on the experiment, an engineer can diagnose the overall condition of the experiment.

5. VISUAL DATA ANALYSIS TOOLS

With the large volume of data that the RAIDS will generate, new ideas about data analysis must be used. Previously, it was not uncommon to store data on tape and forget about it, wasting data. NRL has developed visual data analysis tools (VDATS) that will allow easy performance of data analysis, saving time and allowing more data to be analyzed. The idea is to provide users with a graphical interface that pulls together all of the various aspects of data analysis into one package. These tools are such things as inversion models, simulators, smoothing functions, FFTs, overlays, color patterns, and various graphing and plotting routines. This will allow the user to do those things that he normally does to data, within the confines of the interface. If the user desires something different, the data can be

analyzed interactively using his own methods and customized versions of the software.

In addition to compiling all of the common analysis tools, the VDATs are also easily expandable. This expandability will allow the system to evolve over time to fit nicely into the data analysis process. This expansion would occur when a scientist issues a request to the governing science team and the team approves the enhancement of the system. This process ensures that only those ideas which a majority of users feel would be useful will be added.

After initial data analysis, the user may feel the need to perform certain actions non-interactively, then return later and view the result. While the user is doing the initial data analysis, the VDATs can optionally record all of the actions into a file. This file will then be able to be submitted non-interactively on another data set, allowing easy construction of batch jobs.

As users log on to the VDAT system, they will be placed into the VDAT environment. This environment is a window driven, graphical user interface. The first thing that they will want to do is open a data set and look at the raw, unprocessed data. This is done by choosing 'OPEN' under the 'FILE' selection. After looking at the file, they may want to do some smoothing, change the color table, plot a horizontal or vertical slice of the image, run an FFT on the data, or any number of things. These will all be activated by choosing the appropriate selection from the menu or sub-menus. At any time during the course of this process, the user can select to begin recording the actions for later use by pressing 'LOG' under the 'FILE' selection. As this is chosen, all actions that are performed while in VDATs will be recorded into a file. At the termination of the program, or when the user stops it, recording will stop. The user can then either continue on, or can save the image that he has created and exit the VDATs.

6. APPLICATIONS OF VIDEO RECORDING MEDIA

Another component in the RAIDS system is the standard recording of images and information onto videotape. With the abundance of videotape machines in the world, scientists from NRL will be able to take

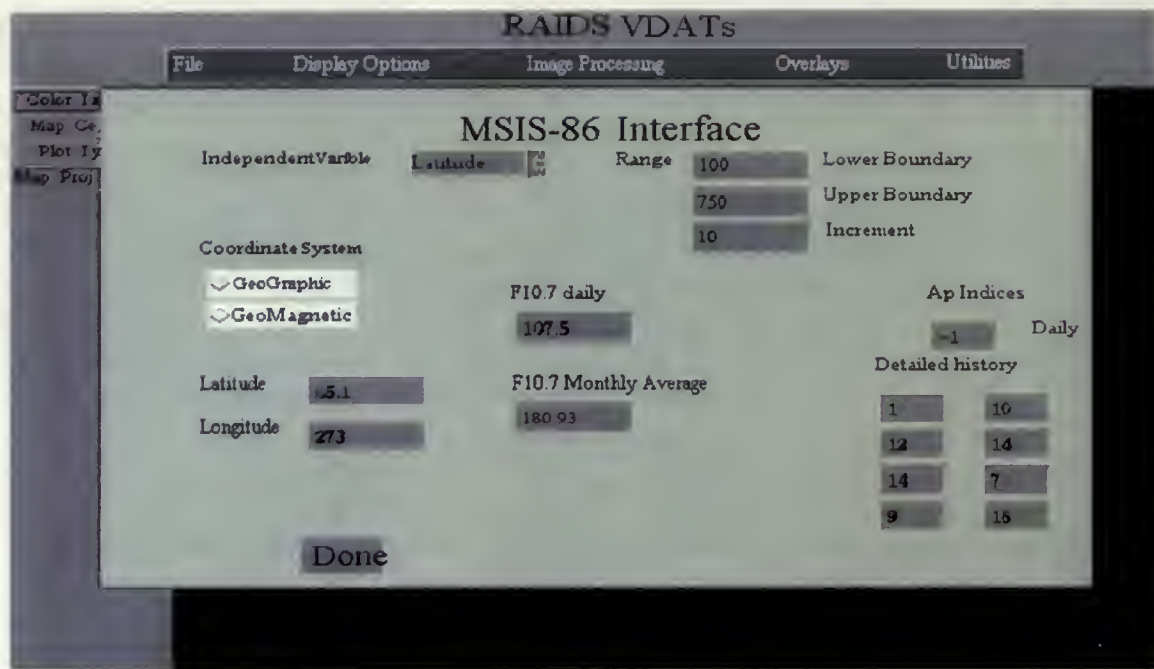


Figure 4. Sample screen showing a model interface. The input parameters to the model are entered into the grey boxes. The background is the main area for the visual data analysis tools.

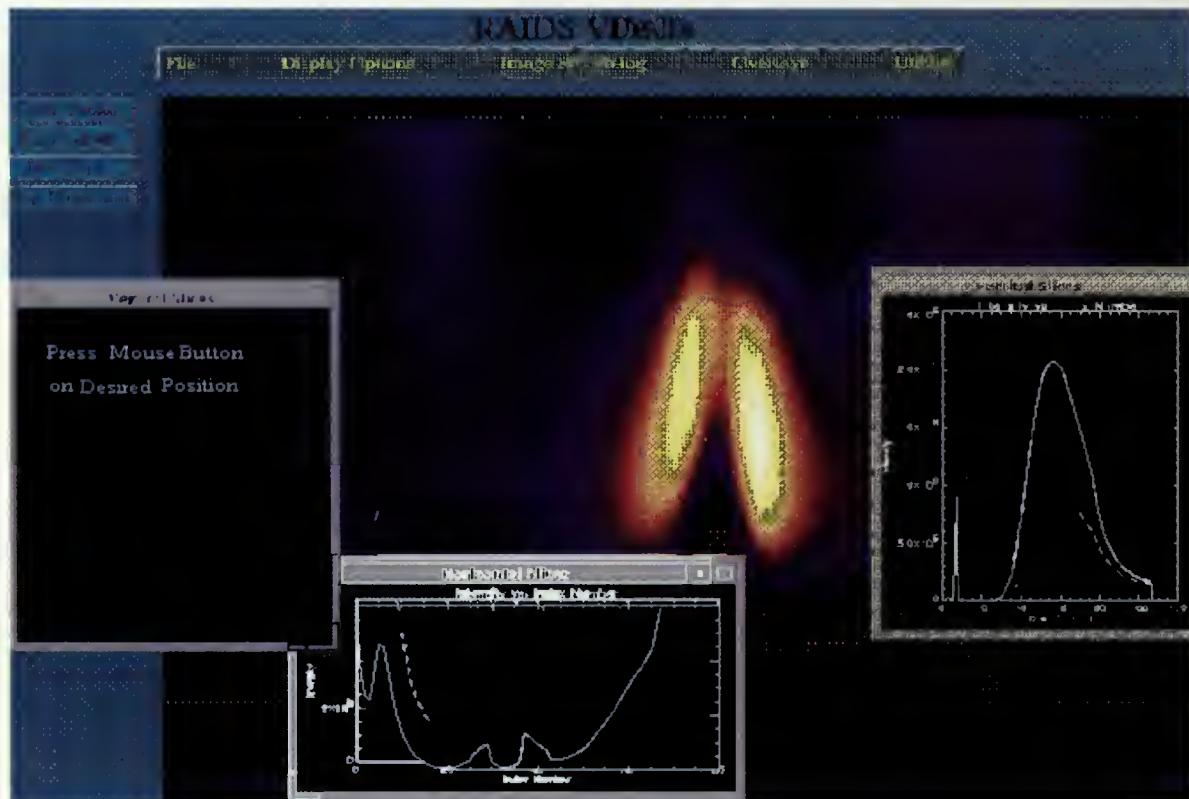


Figure 5. Sample screen showing the image slicing tool. By clicking on a point in the image contained in the main window the row or column that was selected will be plotted in the smaller windows.

data from the laboratory and display the results at professional meetings. The second appealing feature of video recording is the ability to scan or compare a sequence of images for the purpose of identifying interesting, dynamical, or qualitative features. Once a feature is designated for detailed analysis, the original data can be accessed for the frame in which the phenomenon appeared. Screening techniques like this are essential to retrieving maximum scientific information from large data sets.

The process of recording onto videotape is a two step process. The first part is to record the images as they are displayed by the computer on a monitor. This monitor is connected to a "Write Once Read Many (WORM)" optical disc recorder which scans the monitor and can record up to 60 frames per second. The disk can be archived and used as often as necessary to either make new videotapes or for the purposes of previewing data. From the optical disk, the images can be spliced together either in or out of sequence to make a movie. This is then recorded onto video tape for showing at some later time. With the recording onto videotape comes the option of adding or accenting features to make a more useful product. Currently, NRL is acquiring the ability to do video editing of the images, constructing title screens, and even adding audio editing capabilities.

7. SUMMARY

The system described in this paper is a general system that is in place for a general type of instrument. With each new instrument, the same system will be used with modifications in the modular processing software. Through automated processing and visual data analysis, a much larger fraction of the data from the satellites will be able to be used. Instead of creating a new system every time an instrument is created, "front-ends" will be created that will push the data into the standard format.

With the concepts of data visualization and the analysis methods being so new, this system will undergo numerous upgrades. As users begin to gain comfort with it, new ideas will begin. Because the system does not attempt to do everything from the start, while allowing for easy expandability, the system will be able to grow with the analysis community through the years.

Acknowledgments. This project is supported by the Atmospheric Remote Sensing and Assessment Program of the Strategic Environmental Research and Development Program and the Office of Naval Research.

There are many people who have contributed to the data analysis program over the years. The RAIDS engineering visualization was originally developed by the RAIDS engineering team and later rewritten by Sean Parker to take advantage of the graphics now available. The decompression software development has been largely written by Matthew Harrell. The attribute calculation has been written primarily by Dave Hardin and Steve Lockwood. The BDC developed the data base and archival system for RAIDS. Some of the development effort for the graphical data analysis has been supported by the DMSP for the SSULI sensor. In addition, Mark Lundquist has made important contributions to all aspects of the graphical data displays, especially in the area of video presentations and computer graphics.

Satellite Situation Center Data System for Magnetospheric Science Planning

L. AIST-SAGARA^{1,4}, J.F. COOPER^{2,4},
R.E. MCGUIRE¹, R. PARTHASARATHY^{2,4},
M. PEREDO^{3,4}

Critical problems in planning coordinated observation campaigns for magnetospheric science include the need to predict time intervals when one or more observing satellites or ground stations will be connected along magnetic field lines to other observation sites, or when such sites will be located within magnetospheric regions of common interest. The Satellite Situation Center (SSC) was created at the National Space Science Data Center (NSSDC) during the International Magnetospheric Study in the 1970s to address these problems. The SSC Data System has evolved since that era to support potentially complex queries by SSC staff and has now been opened to NASA Science Internet access via the NSSDC On-line Data Information System (NODIS). The SSC software, ephemeris data base, and access modes are described for the Version 2.1 release in 1993.

¹Space Physics Data Facility, NASA Goddard Space Flight Center

²National Space Science Data Center, NASA Goddard Space Flight Center

³ISTP Science Planning and Operations Facility, NASA Goddard Space Flight Center

⁴Technology Applications Group, Hughes STX Corporation

1. INTRODUCTION

The Satellite Situation Center (SSC) is a unit of the National Space Science Data Center (NSSDC) and NSSDC's international counterpart, World Data Center A for Rockets and Satellites (WDC-A-R&S). SSC personnel and software systems support NASA and international space physics activities by maintaining an ephemeris data base for scientific satellites in geocentric or heliocentric orbits, which can be used to plan and support analysis of coordinated science observations by multiple satellites. NSSDC and the Space Physics Data Facility (SPDF), which now provides software development support for SSC, are jointly managed by the Space Science Data Operations Office (SSDOO) at NASA Goddard Space Flight Center with primary contractor support from Hughes STX Corporation.

The SSC was established in the mid-1970s to support and coordinate multi-mission planning for the International Magnetospheric Study (IMS) [Teague et al., 1982]. SSC software resources from the IMS era continued to be used in planning for missions, such as Dynamics Explorer 1 and 2 (DE 12), the International Sun-Earth Explorer series (ISEE 1, 2, and 3), and the Interplanetary Monitoring Platform series (IMP 7 and 8). The SSC supported the ongoing series of Coordinated Data Analysis Workshop studies that began in 1978. In 1986 SSC played a major planning and coordination role during the multi-mission Polar Region and Outer Magnetospheric International Studies (PROMIS) program. Later, SSC similarly supported a joint mission of NASA (U.S.A.), IKI (Russia), and ISAS (Japan) during 1989-90 for coordinated observations with IKI's Active satellite under the aegis of the Inter-Agency Consultative Group (IACG).

Among the projects now supported by SSC are the Solar Terrestrial Energy Program (STEP), NASA's Global Geospace Science (GGS) program, the International Solar Terrestrial Physics (ISTP) program, and the International Heliospheric Study (IHS). SSC data and software have recently been ported to the ISTP/GGS Science Planning and Operations Facility (SPOF) at NASA Goddard Space Flight Center. Precomputed state vector information, predictive or definitive, for the Geotail and WIND spacecraft is being provided by SSC to the IACG-WG3 (Working Group 3) satellite facility, called SPIN, at ISAS. Similar data will be provided to the Joint Science Operations Centre (JSOC) in the United Kingdom Rutherford-Appleton Laboratory for the planned Cluster mission and at the Experimenters Operations Facility (EOF) at GSFC for the SOHO mission. Other space physics data systems, including the new Geospace Environmental Data Display System (GEDDS) at Poker Flat Research Range [Akasofu et al., 1992], which integrates models and ground and satellite observations of the auroral ionosphere for studies of magnetosphere-ionosphere coupling, may be among future users of SSC data.

2. SSC OPERATIONS AND FACILITIES

Updated orbital elements for many satellites are routinely received by SSC on magnetic tape (three per week) from the United States Space Command (USSPACECOM), previously known as NORAD. Elements for satellites of interest to NASA supported missions and international programs are processed by SSC staff into time-ordered Cartesian (X-Y-Z) coordinates, by using the Goddard Trajectory Determination System (GTDS) code maintained at NASA Goddard's Flight Dynamics Division, and stored in a data base within the SSC Software System in NSSDC's Common Data Format (CDF) [Treinish and Gough, 1987]. The Cartesian data points are stored at maximum time resolution of one minute in geocentric inertial coordinates. The orbital element data can be accessed electronically through NASA Science Internet (NSI) at an anonymous FTP directory called ACTIVE. The Cartesian data are available either via the SSC Data System, which is accessed through NSSDC's On-Line Data Information

Service (NODIS), or through a ported version of the data system software on a SUN workstation at the user's home facility.

SSC is supported by software designed to answer common queries about time periods during which one or more satellites and/or ground stations may be in conjunction on the same magnetic field line or in the same magnetospheric region as determined by comparison of spacecraft position and of field line footpoint locations in the ionosphere. The software supports a wide range of user-selected options for internal and external magnetic field models, field line traces, and output coordinate systems. Other capabilities include data listings of satellite coordinates in one of several coordinate systems and identification of magnetospheric region at each point along the satellite track. Calculator functions support conversions of footpoint and satellite locations between different coordinate systems.

The software and hardware capabilities of SSC have evolved over many years since the first generation of SSC programs was written in FORTRAN to run on a MODCOMP IV/25 computer for production of simple reports and data listings. In 1975, an interactive graphics system was added for preconfigured plots of key orbital parameters. In 1985, the software was ported to a MODCOMP Classic II/45 computer after previous updates and additions in 1980 in support of the Dynamics Explorer mission. More recent upgrades have included ports to more powerful computers in the VAX/VMS and SUN/UNIX environments, the addition of new magnetic field models, and the revision of definitions for magnetospheric regions [Parthasarathy et al., 1992].

Whereas user queries were exclusively handled by SSC staff prior to spring 1993, the emergence of the SSC Data System into the NSI network environment now makes possible direct access to SSC software and data by the space science community. The current Version 2.1 release of the SSC Software System, the software component of the data system, has a modern user interface supporting access by X-Windows (X11R5 compatible) and VT-100 compatible terminals, as well as a three-dimensional graphics capability for X/PEX environments supporting the PHIGS (Programmer's Hierarchical Interactive Graphics

Table 1. Time resolution and coverage of satellite data in the SSC Data System (as of August 1993).

Satellite Name	Time Resolution (Seconds)	Definitive/ Predictive (D/P)	Start Coverage Date			End Coverage Date		
			YYYY	DDD	HH	YYYY	DDD	HH
ACTIVE	60	D	1989	272	0.0	1991	277	8.0
AKEBONO	60	D	1989	54	0.0	1993	213	8.0
APEX	60	D	1991	351	0.0	1993	213	8.0
CCE	60	D	1984	230	0.0	1988	1	8.0
CRRES	60	D	1990	206	23.0	1991	277	8.0
	60	P	1991	277	8.0	1992	182	12.2
DE 1	60	D	1981	215	10.6	1991	61	8.0
DE 2	60	D	1981	216	0.0	1983	50	12.3
DMSP-F10	60	D	1990	336	0.0	1993	214	8.0
DMSP-F11	60	D	1991	333	0.0	1993	214	8.0
DMSP-F8	60	D	1990	1	0.0	1993	214	8.0
DMSP-F9	60	D	1990	1	0.0	1992	306	8.0
FAST	60	P	1994	236	0.0	1995	91	0.0
FREJA	60	D	1992	282	0.0	1993	216	8.0
GEOTAIL	720	D	1992	208	0.0	1993	259	8.0
	720	D	1993	259	8.0	1995	1	0.0
GMS 3	60	D	1986	84	0.0	1986	182	0.0
IMP 7	720	D	1972	270	0.0	1978	274	0.0
IMP 8	720	D	1973	303	0.0	1996	1	8.0
INTERBALL (AURORA)	60	P	1994	79	0.0	1996	92	0.0
INTERBALL (TAIL)	720	P	1993	342	0.0	1996	1	0.0
IRM	720	D	1984	256	0.0	1986	242	8.0
ISEE 1	720	D	1977	296	0.0	1987	269	6.6
MOON	1728	P	1991	182	0.0	1993	181	23.5
	1728	P	1993	182	0.0	1996	265	11.0
NOAA 12	60	D	1991	136	0.0	1993	216	8.0
OGO 6	60	D	1969	156	0.0	1970	1	8.0
OHZORA	60	D	1984	51	0.0	1986	102	8.0
POLAR	720	P	1994	141	12.6	1996	143	0.0
RELICT-2	1728	D	1993	349	0.0	1995	1	7.7
SAMPEX	60	D	1992	186	0.0	1993	229	8.0
SCATHA	60	D	1979	36	0.0	1991	1	8.0
SOLAR-A (YOHKOH)	60	D	1991	248	0.0	1993	214	8.0
UARS	60	D	1991	263	0.0	1993	227	8.0
VIKING	60	D	1986	64	0.0	1987	1	0.0
WIND	720	P	1994	75	0.0	1995	32	8.0

System) interface. X/PEX and PHIGS are available from the MIT X Consortium free of any license fees. The user interface was constructed using the JYACC Applications Manager (JAM), which provides a user-friendly menu-driven interface allowing the user to move quickly between different system components and to easily specify run-time parameters and output options.

3. DATA SYSTEM DESCRIPTION

3.1 User Access

The SSC Data System may be accessed over NSI through DECnet or Internet network protocols. The first step is to log onto the NSSDC account `nssdca::nodis` on DECnet or `nodis@nssdca.gsfc.nasa.gov` on Internet. A password is not required in either case. After responding to initial login prompts (type "Y" for the new NODIS interface), the user selects the "Space Physics" option from a menu of major disciplines and then the "Satellite Situation Center" option from the space physics submenu. The user is prompted for a choice of terminal protocols including VT-100, Macintosh with Versaterm Pro (VT/Tek 4105 emulation), Tektronix 4105, SUN Shelltool, and X-term.

3.2 Main Menu

The initial menu of the SSC Data System offers the following options: (1) Query, (2) Locator, (3) Calculator, (4) Database Information, (5) File Processing, (6) Usage Notes, and (7) Exit. Query invokes the Query Processor menu for potentially complex queries regarding field line and/or magnetospheric region conjunctions of multiple satellites and ground stations. Locator invokes the Ephemeris Locator, which provides sequential listings in selectable Earth-centered coordinates of satellite position and region versus time, as well as an additional three-dimensional plot option (SUN-Windows terminals only) for orbital tracks in Geocentric Solar Ecliptic (GSE) coordinates. The Calculator option invokes the Coordinate Calculator for conversion of coordinates between several geocentric systems. Database Information gives the current list of satellites, time resolution, ephemeris

type (usually predictive for pre-launch and definitive for post-launch), and associated date/time coverages. The current list, as of August 1993, is shown in Table 1. The user should note, however, that older ephemeris data sets are now kept off-line, and a special request may be made to bring such data on-line. (Output) File Processing supports the viewing, deletion, printing, and FTP file transfer to the user's home node of the contents of the working directory, including output files from previous runs of the various SSC programs. Usage Notes provide on-line help information. Exit returns the user to the NODIS account.

3.3 Query Processor

The main window for this processor is shown in Figure 1 for the SUN-Windows access option. For selection of one or more satellites and a start/stop time range to meet the specified criteria, a set of conditions on the ephemeris data query can be set up from this window with constraints on the magnetospheric regions (Region Setup) and magnetic field line traces (Trace Setup). The user can specify up to nine sets of conditions for any one query, and a single query may require the current condition, any other condition, or all of the conditions.

The region option enables listing of times during which individual satellites are in one of the selected magnetospheric regions. As shown in Figure 2, the Region Selection Window allows selection of all or only a few regions. In the cases of the dayside/nightside magnetosphere and plasmasphere regions and sub-regions (north, south, cusp, cleft, auroral oval, and polar cap within the magnetosphere), the region choice may exclude one or more sub-regions.

Figure 3 shows an artistic view of the principal region locations in the Earth's magnetosphere. Analytic representations of the region boundaries are set within the software and have been determined from scientific literature and from cumulative experience of, and continuing consultation among, scientific staff of SSC, the ISTEP/GGS SPOF, IACG, STEP, and the scientific community. Pending consensus redefinitions of analytic forms for these boundaries, the region parameters will be updated from those defined for the Version 2.1 SSC software in Parthasarathy et al. [1994].

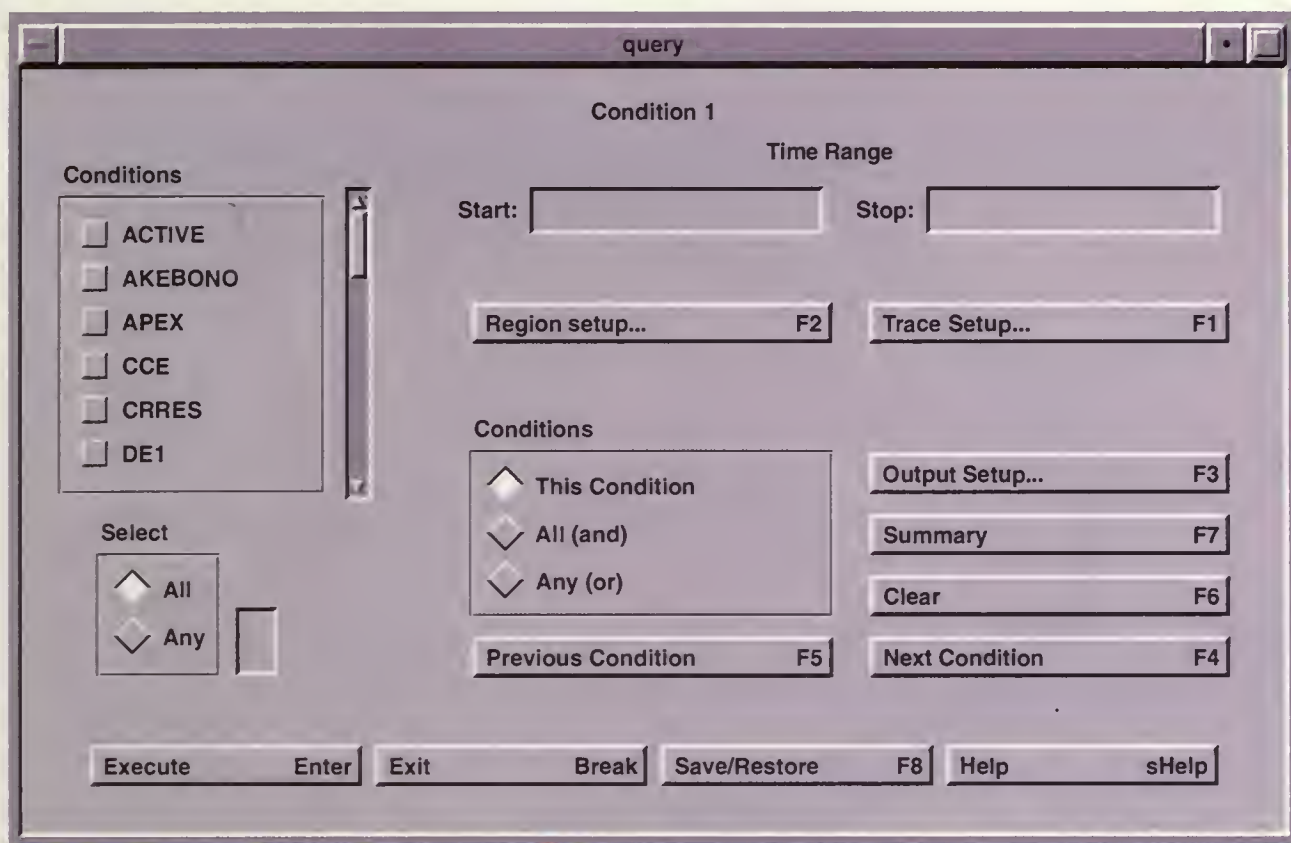


Figure 1. X-Windows version of initial input window for Query Processor in the SSC Data System.

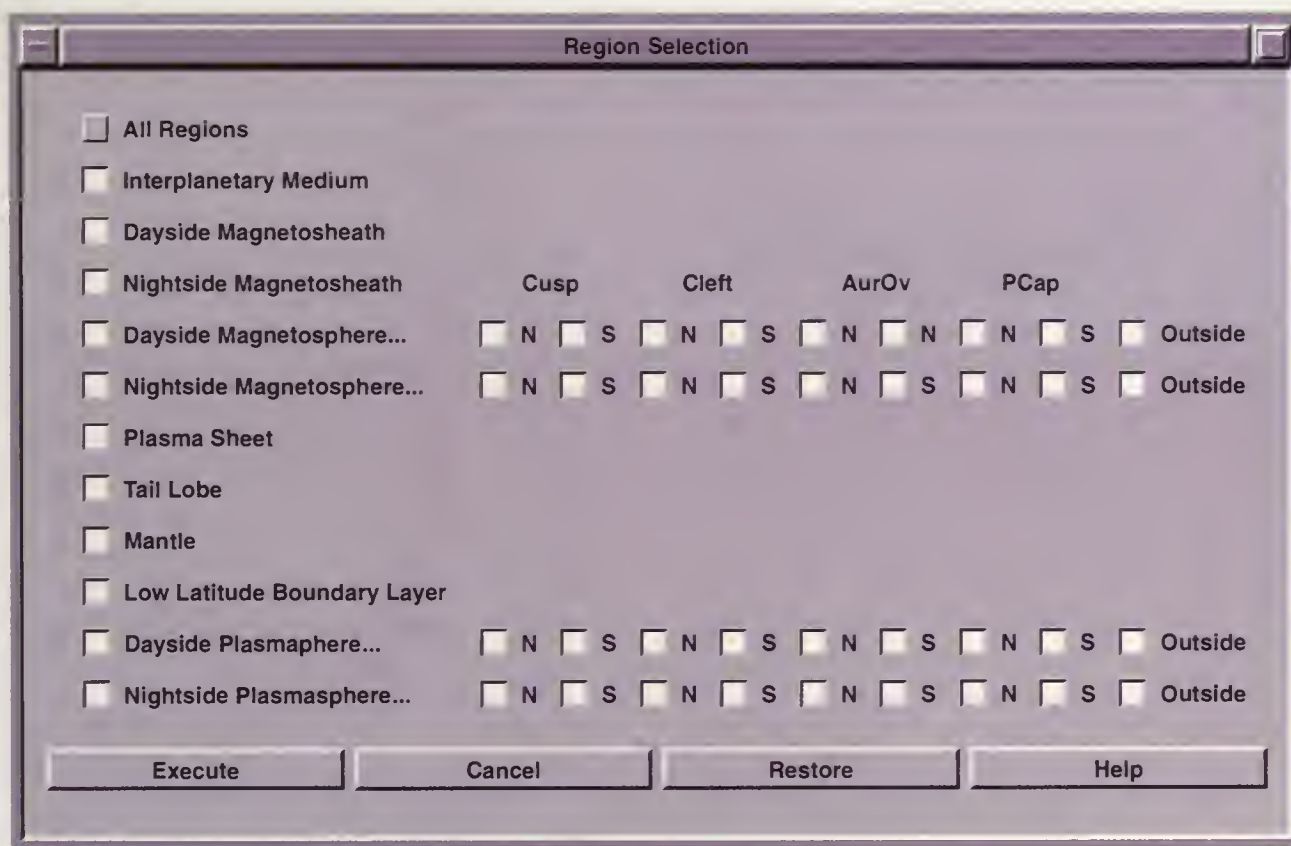


Figure 2. X-Windows version of Region Selection window for Query Processor.




Figure 3. Artistic rendition of regions and boundaries in the terrestrial magnetosphere (not to scale).

The trace option enables field line tracing queries to identify either periods when one or more satellites are on the same magnetic flux tube as a specified lead satellite, or when one or more satellites occupy a field line tracing down near a specified ground station. The flux tube test is done within the specified time range at each orbital point (minimum time separation is 60 seconds) by tracing the local field line of each satellite down to a footpoint at a specified altitude above ground level and by calculating whether that footpoint is within a specified latitude-longitude interval or great circle distance of the specified lead satellite or ground station. Trace options include selectable models for calculation of internal and external magnetic field components at each trace point and selectable direction of trace to determine footpoints in the same, opposite, or both north-south hemispheres.

Magnetic field models now available at NSSDC for the terrestrial magnetosphere have been reviewed by Bilitza [1992]. Internal models supported by Version 2.1 of the SSC Data System include IGRF (Epoch 1965, 1975, 1980), MAGSAT (Epoch 1980), Barraclough (Epoch 1975), and a centered dipole. External models include Tsyganenko 1987 (select long or short tail, warped plasma sheet, Stern variation), Tsyganenko 1989 (select AE or Kp Long), Mead-Fairfield (select superquiet, quiet, disturbed, or super-disturbed), and Olson-Pfizer 1974 (select tilt or no tilt). Model choices will be updated in the future.

3.4 Ephemeris Locator

The Locator supports options to display satellite location data in tabular or graphical format. In the

tabular format the user may request the location in a variety of coordinate systems and with a variety of corresponding parameters shown in the VT-100 window for tabular output options in Figure 4. The following coordinate systems are defined as in Russell [1971]: GEI - Geocentric Equatorial Inertial, GSE - Geocentric Solar Ecliptic, GSM - Geocentric Solar Magnetospheric, SM - Solar Magnetic, GEO - Geographic, and GM - Geomagnetic. The tabular output may be filtered by ranges specified for one or more satellite location parameters.

Panels (a) and (b) of Figure 5 show samples of graphical three-dimensional and two-dimensional output in GSE coordinates, the X axis being directed towards the Sun, the Z axis being directed northward from and perpendicular to the solar ecliptic plane, and tick marks being given along the axis in units of Earth radii. In the two-dimensional plot (Figure 5b), the conical shape is the X-Z GSE

projection of the Sibeck magnetopause boundary [Sibeck et al., 1991] superimposed on the forty day tracks of Geotail, IMP 8, and WIND, based on predictive ephemeris data for 1994. The choice of view direction to give two-dimensional or three-dimensional plots is arbitrary.

3.5 User Support

The online Usage Notes provide some guidance in utilization of the SSC Data System menus, windows, and special keys. A user's guide is available upon request from NSSDC's Coordinated Request and User Support Office (CRUSO). The regular mail address is CRUSO, National Space Science Data Center, Code 633.4, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA. CRUSO can also be reached by phone at (301)286-6695, Fax at (301)286-1771, or by NSI e-mail at request@nssdca.gsfc.nasa.gov or nssdca::request.

Locator
Tabular Output Options

	X	Y	Lt		X	Y	Z	Lt	Ln	Local	
	Z	Ln	Fltr		X	Y	Z	Lt	Ln	time	
GEI	*	-	-		*	*	*	-	-		
GEI/J2000	-	-	-		-	-	-	-	-		
GEO	-	-	-		-	-	-	-	-	-	
GM	*	-	-		-	-	-	*	*	-	
GSE	-	-	-		-	-	-	-	-	-	
GSM	-	-	-		-	-	-	-	-	-	
PACE	-	-	-		-	-	-	-	-	-	
SM	-	-	-		-	-	-	-	-	-	
Trace GEO	-	-	-		-	-	-	-	-	-	
Trace GM	-	-	-		-	-	-	-	-	-	

Specify Add'l output:

	*	Fltr
Regions	-	-
Radial distance	*	-
Distances from:		
Fairfield NS	-	-
bow shock	*	-
magnetopause	-	-
Mag field strengths	*	-
L values	-	-
Invariant latitudes	-	-

Width of output: 78 characters.

Trace Model	F1	Direct the results	Formats&Units	F6
Execute ^T	Clear F9	Summary F5	Back ^E	Help Esc F

RETURN selects the item, TAB moves to the Regions check box, ESC h help

Figure 4. VT-100 version of Tabular Output Options window for Ephemeris Locator.




Figure 5. Graphical outputs from the Ephemeris Locator in GSE coordinates with (left) three-dimensional and (right) two-dimensional views.

4. OTHER RELATED DATA SYSTEMS AT NSSDC

The NODIS account contains several other options which the user may wish to investigate as complementary to the capabilities offered by the SSC Data System. Under the "Multidisciplinary" category in the initial NODIS menu, one can select the NASA Master Directory for information on NASA data systems and mission data sets and the NSSDC Master Catalog for detailed information on spacecraft missions, experiments, and NSSDC-held data sets. In that same category, there are further options for information about access to data sets either directly online through NSSDC's Anonymous FTP account or near-line through the NSSDC Data Archive and Distribution System (NDADS). The ACTIVE directory on the Anonymous FTP account offers frequently updated listings of orbital elements and precalculated field line conjunction events for active scientific spacecraft of interest to the magnetospheric physics community. The

NODIS "Space Physics" category offers direct access to the OMNI data base for near-Earth (e.g., IMP satellite series) interplanetary magnetic field, solar wind plasma, energetic particle, geomagnetic index, and solar activity data, and to a description of the COHO (Coordinated Heliospheric Observations) data base for deep space interplanetary data from missions, such as Pioneer 10 and 11, Voyager 1 and 2, Helios 1 and 2, and Pioneer Venus Orbiter. Heliocentric ephemeris data for such missions, and for major planets, are also available through the ACTIVE and COHO directories on Anonymous. Such data are not currently supported by Version 2.1 of the SSC Data System interface, but may be in future releases.

Acknowledgments. The SSC Software System development has been supported by E. Bugler, D. Hoffman, and L. Sprayregen. F. Ferrier has provided technical assistance in updating of the satellite ephemeris data base for SSC. Past operational support by H.K. Hills and C.M. Wong is also acknowledged.

REFERENCES

- Akasofu, S.I., H.C. Stenbaek-Nielson, T.J. Hallinan, G. Rostoker, and D. N. Baker, A new era in magnetospheric research *EOS*, 73, 305-308, 1992.
- Bilitza, D., Solar-terrestrial models and application software, *Plan. Sp. Sci.*, 40(4), 541-579, 1992.
- Parthasarathy, S., M. Peredo, L. Sprayregen, and R.E. McGuire, Multi-spacecraft mission planning/analysis resources at NSSDC/WDC-A-R&S, in *Solar Terrestrial Energy Program, The initial results from STEP facilities and theory campaigns, Proceedings 1992 STEP Symposium, COSPAR Colloquium Volume 5*, edited by D.N. Baker, V.O. Papitashvili, and M.J. Teague, Pergamon, in press, 1994.
- Russell, C.T., Geophysical coordinate transformations, *Cosmic Electrodynamics*, 2, 184-196, 1971.
- Sibeck, D.G., R.E. Lopez, and E.C. Roelof, Solar wind control of the magnetopause shape, location, and motion, *J. Geophys. Res.*, 96, 5489-5495, 1991.
- Teague, M.J., D.M. Sawyer, and J.I. Vette, The Satellite Situation Center, in *The IMS Source Book. Guide to the International Magnetospheric Study Data Analysis*, edited by C.T. Russell and D.J. Southwood, pp. 112-116, American Geophysical Union, Washington, DC, 1982.
- Treinish, L., and M. Gough, A software package for the data-independent management of multi-dimensional data, *EOS*, 68, 633-635, 1987.

Interactive Visualization of Numerical Simulation Results: A Tool for Mission Planning and Data Analysis

J. BERCHEM, J. RAEDER, R.J. WALKER,
M. ASHOUR-ABDALLA

*Institute of Geophysics and Planetary Physics,
University of California, Los Angeles, CA*

We report on the development of an interactive system for visualizing and analyzing numerical simulation results. This system is based on visualization modules which use the Application Visualization System (AVS) and the NCAR graphics packages. Examples from recent simulations are presented to illustrate how these modules can be used for displaying and manipulating simulation results to facilitate their comparison with phenomenological model results and observations.

1. INTRODUCTION

With the advent of multispacecraft missions, such as the NASA Global Geospace Mission [e.g., Russell, 1994], it has become obvious that theoretical models are needed, not only to interpret local spacecraft observations, but also to link single point measurements. To these ends, mission theorists have been developing models that address space plasma phenomena on local and global scales, involve nonlinear processes, and use self-consistent approaches. Because of the complexity of these models, it has become standard to use numerical simulation techniques.

In the past, numerical simulations have generally been used to provide an indepth physical understanding of phenomena that had first been interpreted in terms of simple models and linear theory. As such, simulations came at the end of the intellectual process of interpreting spacecraft measurements, rather than at the beginning. From its inception, the International Solar Terrestrial Physics (ISTP) program coordinators recognized that it was necessary to develop the ability to use simulations at an earlier stage of the mission for use in the planning process. An additional challenge was to increase the realism of the simulation models so their results could provide quantitative predictions that could be effectively compared to specific spacecraft observations.

In response to these challenges, ISTP theoretical teams have been producing more and more sophisticated simulations by refining algorithms, as well as improving the temporal and spatial resolution. The large volume of information produced by these numerical simulations has presented new problems in data management and analysis. The recent advances in both hardware and software

graphics technology have not only considerably helped to alleviate these problems but also have opened new and exciting possibilities. Visualization techniques have become invaluable, not only for displaying numerical simulation results, but more importantly for analyzing the implications of these results.

2. NUMERICAL SIMULATION DATA SETS

2.1 Simulation Techniques

Since the 1960s, numerical simulations of plasmas have been used extensively to study a vast array of physical problems, ranging from laboratory plasmas to interstellar plasmas. Numerous simulation techniques have been developed which depend on both the temporal and the spatial scales of the plasma phenomena investigated.

One can distinguish two categories of approaches frequently used to model the time-dependent, collisionless plasmas normally encountered in space. The first class is made up of "kinetic" simulations. These simulations are the equivalent of following the evolution of the Vlasov equation together with Maxwell's equations by calculating the orbits of a large number of finite-sized particles in their self-consistent electromagnetic fields. There are many varieties of particle codes, including electrostatic, magnetostatic, and electromagnetic models [e.g., Hockney and Eastwood, 1981; Birdsall and Langdon, 1985]. Since particle codes give a full description of the local plasma physics, they have been widely used to investigate the plasma instabilities generated by wave-particle interactions that occur in many regions of the magnetosphere. There are many additional particle simulation approaches, such as methods which implicitly integrate particle motion [e.g., Brackbill and Forslund, 1982], methods which retain only particles' slow drifts by averaging over their more rapid gyromotion [e.g., Lee, 1987], and "hybrid" methods where one of the species (usually the electrons) is described as a fluid [e.g., Winske, 1985].

At the other end of the spectrum are the fluid simulation models. Although collisions are infrequent, correlation times are small enough to use magnetohydrodynamic (MHD) simulations to model large scale processes. These codes self-consistently solve a closed system of Maxwell's

equations and ideal fluid equations as an initial and boundary-value problem. They have been used quite successfully to model local phenomena, as well as the global interaction of the solar wind with Earth's magnetosphere [e.g., Ogino et al., 1986; Fedder and Lyon, 1987]. Fluid models based on higher moment transport equations are also being used to describe non-equilibrium and multifluid processes, which are particularly important in ionosphere-magnetosphere coupling and cometary mass loading [e.g., Schunk, 1977; Gombosi and Körösmeszey, 1989].

MHD simulations have inherent limitations in rendering kinetic effects, such as gradient and curvature drifts. To evaluate the importance of kinetic effects in a large system other approaches have been used. One of them is based on the numerical integration of particle trajectories in a given set of electromagnetic fields and involves calculating physical quantities by deriving moments obtained from the particles' distribution functions. Although the computed distributions are not necessarily consistent with the electromagnetic fields, these "large-scale kinetic" simulations have a distinct advantage because they can be used with analytical, as well as data-constrained models. They have been very useful in investigating the macroscopic signatures of kinetic processes, such as stochastic scattering in the magnetotail [e.g., Ashour-Abdalla et al., 1993].

The simulation techniques discussed above have been implemented in many different ways, depending on the specific area of magnetospheric physics to which each was applied. Details of the simulation algorithms currently being used in space plasma physics can be found in the proceedings from the International School for Space Simulations [Matsumoto, 1982; Ashour-Abdalla and Dutton, 1985; Lembege and Eastwood, 1988; Matsumoto and Omura, 1991].

2.2 Simulation Diagnostics

A great variety of diagnostics is used to analyze simulation results. Most of the time these diagnostics involve multi-dimensional space and are presented graphically. Although a large fraction of them are generic, such as contour plots of electromagnetic fields, densities, etc., some diagnostics are

more specific to the technique used to carry out the simulation. For example, diagnostics of most kinetic simulations include phase space representations of particle distributions and wave spectral analyses to identify sources of free energy and wave-particle interactions. In fluid-type simulations, stronger emphasis is given to diagnosing topological properties of flows and electromagnetic fields by tracing, for example, stream lines and magnetic field lines.

The main difficulty is not in calculating these diagnostics, but rather in implementing their computation with as little overhead as possible on the time-step cycle. This depends of course on the performance and capabilities of the computers and peripherals used, in particular the input/output speeds and the internal and external mass-storage capacities. In early numerical simulations, digital output was held to a minimum. Almost all diagnostics were embedded in the code because of slow output speeds and the rapid saturation of storage space. This limited the analysis that could be performed on the simulation results since diagnostics had to be implemented before the run. Recent progress in computer technology has considerably reduced these bottlenecks. Today most of the diagnostic routines can be dissociated from the simulations to speed up the execution time. Large data sets can be extracted from the computation and

stored to be post-processed either as the computation is executed or later on. As we will see later, the average size of an output file from current three-dimensional MHD simulations is of the order of 16 MB per snapshot in time. Files of this size can rapidly add up to massive data sets (1 GB per simulated hour when following the evolution of the magnetosphere with a time resolution of one minute).

2.3 Display of Simulation Results

Despite the massive processing capabilities of modern computers, space research is still fundamentally dependent on the investigator's ability to visually recognize patterns in the data. Experimentalists have invested a tremendous amount of effort in creating efficient displays for the meaningful representation of large data sets. It has long been recognized that to be effective, simulation results should, to the greatest extent possible, be presented in formats that parallel the data displays used for spacecraft measurements [e.g., Ashour-Abdalla et al., 1994] to help assess whether theoretical predictions are consistent with spacecraft measurements.

We have already started to produce simulation results in an experimental type of format. Figure 1 shows results of a two-dimensional electrostatic simulation of an ion beam-driven system which was




Figure 1. Results of a two-dimensional electrostatic simulation showing wave diagnostics displayed in formats similar to those used to plot measurements from wave experiments onboard spacecraft. The right panel exhibits frequency-wave number dispersion relations color coded according to wave power.

designed to model the generation of broadband electrostatic noise (BEN) observed in the plasma sheet boundary layer [Gurnett et al., 1976]. Although BEN is thought to result from the electron acoustic and ion-ion instabilities generated by field-aligned beams [e.g., Schriver and Ashour-Abdalla, 1990], we display here only the part of the wave analysis that indicates that electron harmonic cyclotron (ECH) waves can be simultaneously enhanced during the process [Berchem et al., 1991].

The left and center panels of Figure 1 show the results from the simulation displayed in formats similar to those used to plot Sweep Frequency Receiver (SFR) spectrograms and power spectra obtained from wave experiments onboard spacecraft. The right panel exhibits frequency-wave number dispersion relations color-coded according to wave power. Although wave instruments do not obtain a direct measurement of wave numbers, this plot shows how we can exploit the extra information produced by simulations. By comparing this diagram to linear wave theory we can clearly recognize the classical dispersion relation expected for ECH waves and positively identify the peaks observed in the spectrogram and the power spectrum.

Recent advances in both hardware and software have provided new possibilities for analyzing simulation results other than simply duplicating experimentalists' formats. As we will see in the next section, it is now relatively straightforward to design and implement interactive visualization environments where simulation results can be analyzed in the context of observations to realize synergistically the goals of mission-oriented theory.

3. THE UCLA MISSION-ORIENTED VISUALIZATION SYSTEM

3.1 *The Development of the Mission-Oriented Visualization System*

The UCLA Mission-Oriented Visualization System is an interactive post-processor of simulation results that uses the commercial Application Visualization System (AVS) software for three-dimensional volume rendering and the NCAR graphics package for two-dimensional plots. AVS is a modular programming environment based on a data-flow

network architecture. Originally developed in 1989 by the Stardent Computer Company, this software allows users to build visualization networks by simply connecting individual modules using mouse-driven point-and-click operations. These modules include a variety of data-input, filtering, mapping, and rendering output routines that provide a comprehensive set of visualization tools. In addition to providing an intuitive and easy-to-use interface for configuring, modifying, and saving networks, the AVS architecture lets the user create custom modules to control any aspect of AVS. This capability provides the true power of the AVS package. Being able to write and incorporate modules has allowed us to design our own application system to take full advantage of the AVS generic visualization environment while meeting our needs in displaying and analyzing space plasma simulation data sets. This task can be greatly facilitated by using the AVS module generator that creates skeletons into which users can insert specific algorithms.

Our system development took place on two levels. Most of the standard modules (aside from the basic tools such as color maps, slicers, and viewers, etc.) are either too general or too limited for developing a specific application, although they are very useful at the prototyping stage. The first part of the development was to design and implement new modules that were needed to visualize and analyze our simulation results in a magnetospheric context. In the next section, we will give a few examples of modules that were designed for analyzing three-dimensional global MHD simulation results. The second part of the system development was focused on improving the user interface. Although the AVS interface is intuitive and easy to learn, it can be tedious and frustrating to assemble complicated networks or just to remember the appropriate sequential combinations of operations required to perform certain tasks. Most of the time, predefined menus are easier to work with as they provide straightforward access to the most frequently used analysis routines. This approach also allows users who know only minimal AVS syntax to perform some analyses quickly. Nevertheless, we designed the interface such that primary networks remain accessible to users who are more familiar with AVS in order to let them incorporate other elements into their analysis.

3.2 Global MHD Simulations

The results that we use to illustrate our visualization system were obtained from three-dimensional MHD simulations of the solar wind interaction with Earth's magnetosphere. These global simulations were carried out specifically to show how they could help locate the four spacecraft of the CLUSTER mission (launch planned for end of 1995) with respect to the various magnetospheric boundaries [Berchem et al., 1993]. Our purpose at this stage was not to conduct a systematic study, but to provide practical examples of what could be done for mission planning. We imposed two simple upstream boundary conditions (northward and southward) on the incoming interplanetary magnetic field (IMF) and described the ionospheric boundary of the system by using a conductivity tensor obtained from an analytical model [Rasmussen and Schunk, 1987]. We focused on a single period of the year (the vernal equinox) by specifying in the model the actual orientation of the geographic and magnetic dipole axes for that period.

The algorithm used in this simulation solves one-fluid ideal MHD equations with an explicit

conservative predictor-corrector time-stepping scheme and hybridized numerical fluxes for fourth order spatial finite differencing. The computational mesh is rectangular (x, y, z) but nonuniform, and contains 10^6 grid points. For the results shown here, we designed a special grid to obtain maximum resolution ($0.4 R_E$) in the cusp regions. The dimensions of the simulation box are $32 R_E$ in the sunward direction, $80 R_E$ along the tail and $\pm 35 R_E$ in each transverse direction. Simulations using larger system sizes (up to $400 R_E$ tailward) have also been carried out for distant tail studies in support of the ISTP GEOTAIL mission [e.g., Raeder et al., 1993]. Basic information on the initialization procedures and boundary conditions of the simulations reported here can be found in Berchem et al. [1993].

3.3 Visualization of Global Magnetospheric Configurations

Figure 2 shows simulation results obtained after one hour of real time has elapsed during which a steady northward IMF boundary condition has been maintained to convect the unphysical initial state outside of the simulation system. Only a small part




Figure 2. Perspective view of two-dimensional color-coded contours of the logarithm of plasma pressure obtained from three-dimensional global MHD simulation for northward IMF.

($100 \times 50 \times 50 R_E$) of the simulation system is represented by a three-dimensional perspective view of two-dimensional color-coded contours of the logarithm of plasma pressure. These cuts are obtained by using standard AVS orthogonal slicers. One cut is taken along the meridian x - z (GSE) plane, whereas the two others are taken orthogonally to that plane at $x = 30$ and $75 R_E$. The strong skewing of the magnetotail results from the tilt of the magnetic dipole axis.

We traced two-dimensional field lines obtained from the magnetic field vectors projected onto the meridian contours of the plasma pressure. Although they are not strictly magnetic field lines because of the tilted dipole axis used here, these traces indicate clearly that reconnection with the solar wind field is occurring in the nightside high-latitude region, as is expected during periods of northward IMF. Since the tilt angle is small, the lines also help in visualizing the different magnetospheric regions. Starting at the left with a low pressure solar wind region (≈ 5 pPa; deep blue), we observe successively the bow shock, marked by a narrow layer of increasing pressure (yellow), a high pressure magnetosheath region (≈ 600 pPa; red), another transition region of intermediate pressure corresponding to the magnetopause and the low-latitude boundary layer (orange and yellow), the plasma mantle (green and turquoise) where the pressure starts to decrease towards its minimum magnetospheric value in the lobe (deep blue), then the plasma sheet boundary layer where it rises again (turquoise) and finally the central plasma sheet with a pressure comparable to its mantle value (green). Note that the deep blue $3.7 R_E$ circular region with its center at Earth delimits the inner boundary of the simulation model.

When analyzing a snapshot of three-dimensional MHD simulation results, it is very important to work from a global understanding of the state in which the simulated magnetosphere has evolved before moving on to a deeper analysis. Identifying regions where magnetic field energy is converted into plasma energy or the reverse, is one of the fundamental parts of this process. Tracing magnetic field lines is of course essential in localizing merging processes. Nevertheless, other diagnostics, such as plotting the electromagnetic quantity ($\mathbf{E} \cdot \mathbf{J}$) or the ratio of the plasma pressure over the mag-

netic field pressure (plasma beta), are also very instructive. We implemented a module that can extract and calculate these physical quantities, which can then be displayed by using standard AVS modules or fed to other modules for further manipulation. From the five primary fields routinely extracted from the computational grid (plasma pressure, density, and velocity; magnetic field and current density), more than 30 quantities are derived. Some of these quantities are just vector components and magnitudes of primary quantities or straightforward plasma parameters and are computed interactively. Other quantities, such as the magnetic field topology parameter (this parameter indicates whether a grid point is located in an open, a reconnected, an Earth-closed, or a loop-closed magnetic field line), are more involved and need to be either generated during the simulation run in addition to the regular grid parameters or post-processed.

3.4 Visualization of Three-Dimensional Local Structures

Although they require further refinement, global MHD simulations are now sufficiently sophisticated that they can reproduce small scale features. The substantial increase in spatial resolution obtained during the past few years has been made possible not only by improving algorithms but also using more and more powerful computers such as massively parallel machines. The simulations shown here, for example, were run on the San Diego Supercomputer Center (SDSC) Intel iPSC/860 and Paragon machines. These new capabilities prompted us to design modules to facilitate the visualization of local processes.

One of these modules is our interactive field line tracer. As we mentioned before, tracing magnetic field lines is very important for understanding the global topology of the magnetosphere; it can be very useful for local studies as well. Although a certain number of field (or stream) line tracers is available, they are not flexible enough to carry out such local analyses. The three-dimensional topology of magnetospheric processes can be fairly complex. For example, the formation of magnetic flux ropes and closed loops can lead to endless field lines that are sources of

problems in terms both of computational speed and memory space. Selecting the field lines to be plotted presents another difficulty. In most cases, one attempts to isolate and follow bundles of field lines rather than using large scale systematic displays which result in intricate and unintelligible plots.

Our field line tracer module uses a mouse-driven point-and-click technique to determine interactively the starting points of the individual or group of field lines to be traced. In addition, the integration step, the length of the field line and the number of segments to be displayed can be adjusted. Field lines are automatically color coded to indicate whether they are open, closed, etc. This module is thus particularly useful for visualizing magnetic field lines, but it can be used as well to study the topology of any of the vector fields calculated or derived from the simulations. We have already shown (Figure 2) an example of an application in which we used the vector field obtained by projecting the magnetic field onto a plane. Figure 3 shows another example where we have used the field line tracer to obtain a three-dimensional rendering of the dayside magnetopause Chapman-Ferraro current system.

In Figure 4, we display an example of the field line tracer used in conjunction with an isosurface representation to investigate the three-dimensional structure of the magnetic cusps for northward IMF. The isosurface of the plasma pressure representing the dayside magnetospheric boundary is viewed from the sun (Panel a) and from above the northern cusp (Panel b). We have traced only the magnetic field lines originating from the northern cusp region to clarify the representation. Careful examination of these views reveals that two large funnel-like structures are formed in each cusp region. One of the funnels is located at lower latitudes and is populated only by inner closed field lines, which are shown in black. The second funnel, found at higher latitudes, is composed of a mixture of closed (blue) and open (yellow) field lines. The open field lines drape over the closed ones; their open ends are located in the southern hemisphere. These field lines are newly reconnected field lines which are convected along the magnetospheric flanks. Analysis of flow patterns and the ionospheric electrostatic potential (not shown here) suggests that this type of structure in the plasma pressure appears because the dipole tilt destroys the symmetry of the convection patterns [Berchem et al., 1993].




Figure 3. Rendering of the Chapman Ferraro current system. Panel (a) shows the contours of plasma pressure in the x - z meridian plane that were used to select points (shown by the green spheres) located at the magnetopause, whereas Panel (b) shows traces of the current lines passing through these points.

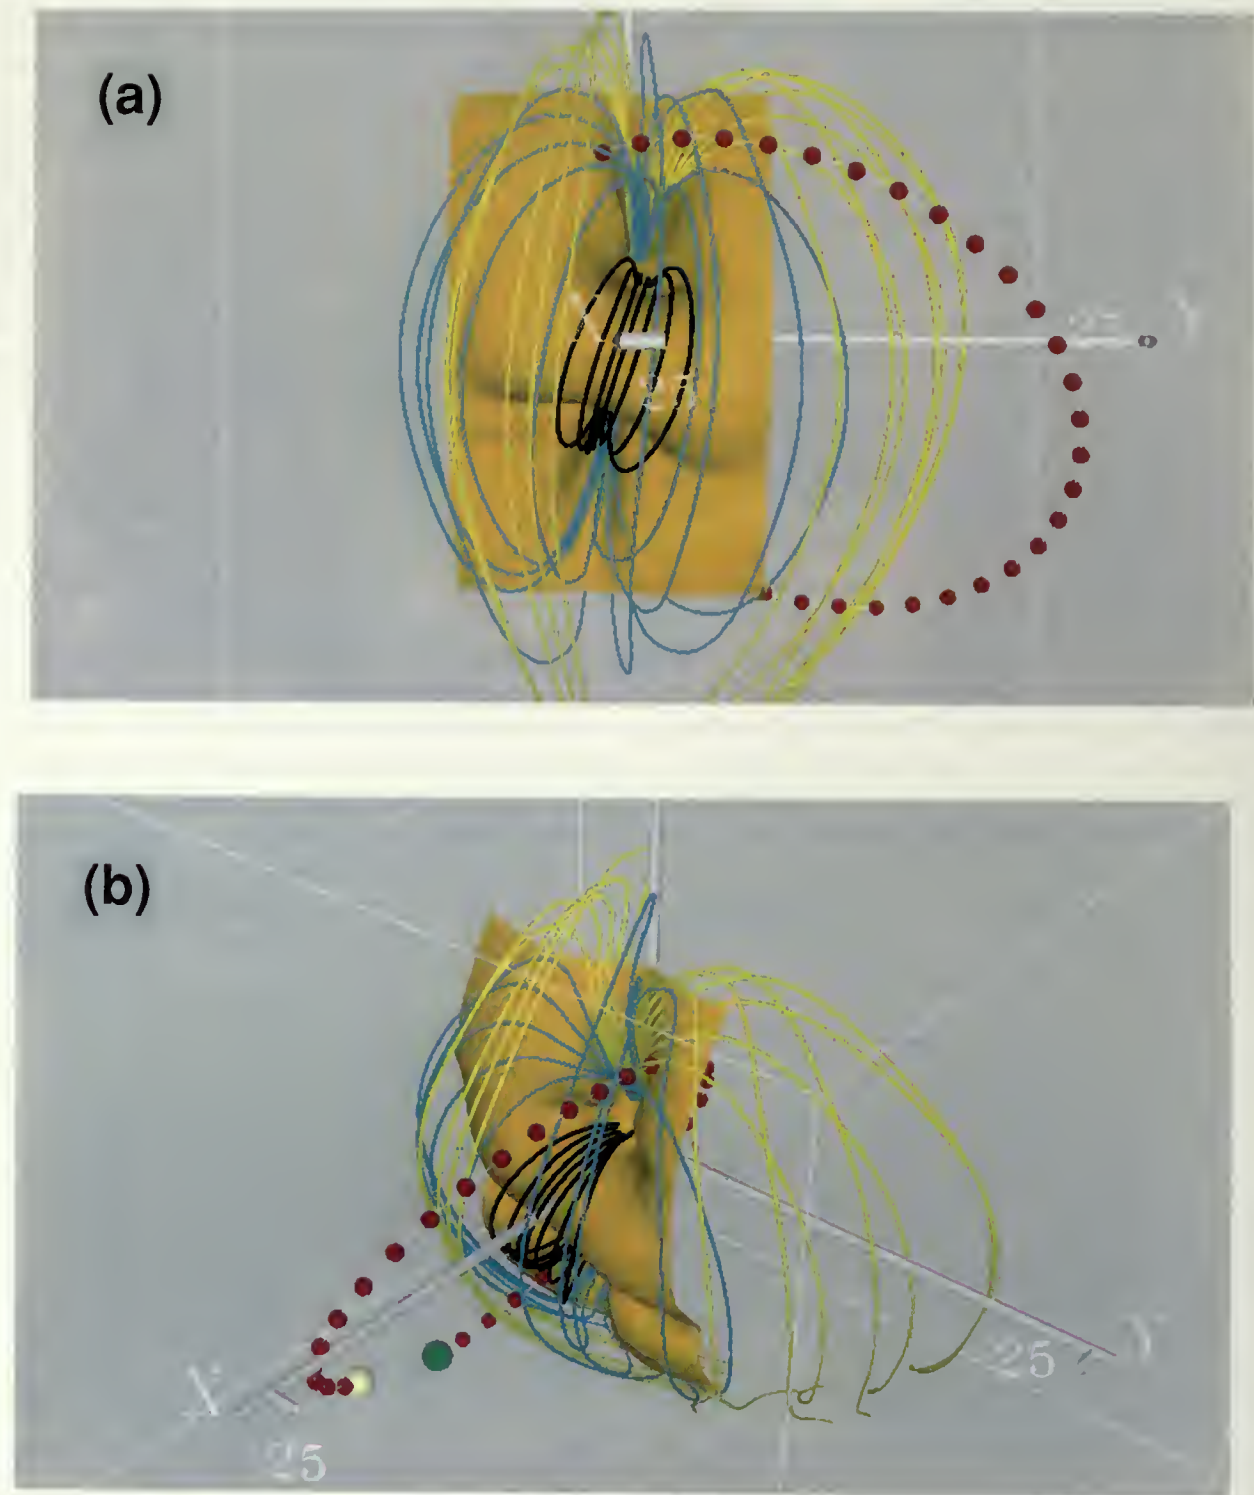


Figure 4. The isosurface of the plasma pressure representing the dayside magnetospheric boundary is viewed from the sun (Panel a) and from above the northern cusp (Panel b). We have traced only the magnetic field lines originating from the northern cusp region to clarify the representation.

3.5 Data Closure Through Visualization

As we have seen above, the visualization of both global and local magnetospheric structures is a very important tool for studying the evolution of the magnetosphere. However, except for being intellectually satisfying, a qualitative comparison between these structures is not sufficient to realize the data closure step necessary to improve the model. Comparing actual satellite measurements with simulated data streams from virtual spacecraft in fact may provide the most direct way of determining whether the phenomena represented by the simulation are the best elements to use for interpreting the spacecraft observations. To add this important capability to our system, we designed a special module that displays real or fictitious spacecraft trajectories and plots the simulation results along them. Input to this module can be in the form of an orbit file or free-hand selection using a mouse-driven point-and-click technique. In addition, we can compute interactively the values from the data-constrained Tsyganenko magnetic field models along the orbit and superimpose them onto the plots. This module uses the NCAR graphics package because it offers more versatility than AVS does in displaying two-dimensional plots. In addition to being displayed in a screen window, the data stream can be written to PostScript and ASCII files and then printed or further analyzed by traditional time series analysis programs.

Figures 5 and 6 show an example of the output obtained for one of the inclined ($8 \times 22 R_E$) CLUSTER trajectories that has been considered in mission planning. In Figure 5 we have superimposed the spacecraft trajectory over the contours of the logarithm of the plasma pressure in the orbital plane. Only the front part ($25 \times 25 \times 25 R_E$) of the simulation system is displayed here, viewed from the duskside at about 20:00 LT. The spacecraft trajectory is shown by small red spheres linked by a white line. The green and yellow spheres indicate the start and end of the trajectory respectively. Simulated data streams are shown in Figure 6 for both southward (left) and northward (right) IMF. From the top to the bottom, we plotted the three components (GSE) and magnitudes of the magnetic field (nT), plasma flow velocity (km/s), density (cm^{-3}), pressure (pPa), and plasma beta. To assist




Figure 5. Spacecraft trajectory superimposed over the contours of plasma pressure in the orbital plane.




Figure 6. Simulated data streams for both southward (left) and northward (right) IMF.

in making comparisons, we superimposed (dashed lines) the magnetic field obtained from the Tsyganenko [1989] model for a time at the center of the time period used to calculate the spacecraft trajectory and option 2 ($K_p = 1-, 1, 1+$). Note that

the values from the magnetic field model [Tsyganenko, 1989] are not valid beyond the magneto-pause. Locations of the spacecraft are given at the bottom of the figure; R is the geocentric distance (R_E), LT the local time and LAT the latitude (degrees); all are calculated in GSM coordinates.

Although the resolution of the plots shown here is relatively low because we had to accommodate the large variations observed along the entire trajectory, we can easily follow the virtual spacecraft through the different magnetospheric regions and identify their boundaries. Special ports to the module have been designed to superimpose spacecraft data when they become available. Since the period of steady solar wind modeled here is probably too long, we did not attempt to make any specific comparisons with existing data. However, we anticipate that this system will be even more valuable when we have observational data for which solar wind conditions are known, as these can be used as input parameters for the global simulations. This will allow us to move to the next step, that of thoroughly testing the predictive capabilities of our models.

4. SUMMARY

We reviewed here examples illustrating a few of the capabilities of the interactive visualization system that we developed to analyze the results of mission-oriented simulations. These examples show that recent advances in hardware and software graphics technology have opened new and exciting possibilities. As in many other areas, visualization techniques have become an invaluable tool for space physicists. They can be used to interact with the results of very complicated simulation models in order to grasp the fundamental physics involved and compare their results to observations. With these new interactive capabilities, visualization has also evolved from being the end product of the physical analysis (the "pretty picture") to becoming an integral part of the intellectual process of exploring, testing, and communicating simulation results. We have seen through the examples shown here that simulation and visualization techniques have reached the stage where together they can provide a

very powerful tool to interpret space plasma observations and realize the synergistic theory-data closure necessary to improve further our theoretical models. We anticipate that this tool will also be very useful for planning future missions, as well as in defining the best strategy for coordinated theory-data campaigns.

Acknowledgments. We thank T. Ebrahimian for his help in developing the visualization software. Simulations were performed on the iPSC/860 and the Paragon machines at the San Diego Supercomputer Center. Funding was provided by NASA ISTP Grant NAG5-1100. IGPP UCLA publication #3973.

REFERENCES

- Ashour-Abdalla, M., and D.A. Dutton (Eds.), *Space Plasma Simulations*, 578 pp., D. Reidel, Boston, MA, 1985.
- Ashour-Abdalla, M., J. Berchem, F.V. Coroniti, J. Raeder, D. Schriver, and R.J. Walker, Mission-Oriented theory for ISTP, *Space Sci. Rev.*, in press 1994.
- Ashour-Abdalla, M., J. Berchem, J. Büchner, and L.M. Zelenyi, Shaping of the magnetotail from the mantle: Global and local structuring, *J. Geophys. Res.*, 98, 5651-5676, 1993.
- Berchem, J., J. Raeder, and M. Ashour-Abdalla, Mission-oriented simulations of small and large scale structures in the dayside magnetosphere, in *Spatio-Temporal Analysis for Resolving Plasma Turbulence*, p. 393-408, ESA WPP-047, 1993.
- Berchem, J., D. Schriver, and M. Ashour-Abdalla, Simultaneous excitation of broadband electrostatic noise and electron cyclotron waves in the plasma sheet, *Geophys. Res. Lett.*, 18, 729, 1991.
- Birdsall, C.K., and A.B. Langdon, *Plasma Physics via Computer Simulation*, 479 pp., McGraw-Hill, New York, NY, 1985.
- Brackbill, J.U., and D. Forslund, An implicit method for electromagnetic plasma-simulation algorithms, *J. Comput. Phys.*, 46, 271-307, 1982.

- Fedder, J.A., and J.G. Lyon, The solar wind-magnetosphere-ionosphere current voltage relationship, *Geophys. Res. Lett.*, **14**, 880-883, 1987.
- Gombosi, T.I. and A. Körösmezey, Modeling of the cometary nucleus-coma interface region, *Adv. Space Res.*, **9**, 41, 1989.
- Gurnett, D.A., L.A. Frank, and R.P. Lepping, Plasma waves in the distant magnetotail, *J. Geophys. Res.*, **81**, 6059, 1976.
- Hockney, R.W., and J.W. Eastwood, *Computer Simulation Using Particles*, 540 pp., Adam Hilger, New York, NY, 1988.
- Lee, W.W., Gyrokinetic particle simulation model, *J. Comput. Phys.*, **72**, 243-269, 1987.
- Lembege, B., and T. Eastwood (Eds.), *Numerical Simulation of Space Plasmas*, 274 pp., Elsevier, Amsterdam, 1988.
- Matsumoto, H. (Ed.), *The First International School for Space Simulations (ISSS)*, 232 pp., Kyoto Univ., Kyoto, Japan, 1982.
- Matsumoto, H., and Y. Omura (Eds.), *The Fourth International School for Space Simulations (ISSS)*, 342 pp., Kyoto Univ., Kyoto, Japan, 1991.
- Ogino, T., R.J. Walker, and M. Ashour-Abdalla, An MHD simulation of the effects of the interplanetary magnetic field B_y component on the interaction of the solar wind with the Earth's magnetosphere during southward interplanetary magnetic field, *J. Geophys. Res.*, **91**, 10,029-10,045, 1986.
- Raeder, J., R.J. Walker, and M. Ashour-Abdalla, The structure of the distant geomagnetic tail during long periods of northward IMF, *Geophys. Res. Lett.*, submitted 1993.
- Rasmussen, C.E., and R.W. Schunk, Ionospheric convection driven by NBZ currents, *J. Geophys. Res.*, **92**, 4491-4504, 1987.
- Russell, C.T., (Ed.), *Global Geospace Mission*, Kluwer Academic Publishers, Dordrecht, The Netherlands, submitted 1994.
- Schrifer, D. and M. Ashour-Abdalla, Cold plasma heating in the plasma sheet boundary layer: theory and simulations, *J. Geophys. Res.*, **95**, 3987-4005, 1990.
- Schunk, R.W., Mathematical structure of transport equations for multispecies flows, *Rev. Geophys. Space Phys.*, **15**, 429, 1977.
- Tsyganenko, N.A., A magnetospheric magnetic field model with a warped tail current sheet, *Planet. Space Sci.*, **37**, 5, 1989.
- Winske, D., Hybrid simulation codes with application to shock and upstream waves, *Space Sci. Rev.*, **42**, 53-66, 1985.

Chapter III

**LOWER
ATMOSPHERIC
AND
EARTH
SCIENCE
APPLICATIONS**

Using the Application Visualization System To View Haloe Three-Dimensional Satellite Data

MINGZHAO LUO¹, ALLEN V. R. SCHIANO²,
JAMES M. RUSSELL III³, LARRY L. GORDLEY⁴,
KENNETH A. STONE⁴, RALPH J. CICERONE¹

The Application Visualization System (AVS) is used to view a three-dimensional data field containing the volume mixing ratios of a chemical species in the middle atmosphere obtained by the Halogen Occultation Experiment (HALOE) aboard the Upper Atmosphere Research Satellite (UARS). Since launch in September 1991, HALOE has been collecting data on approximately 30 sunrise/sunset events in two narrow latitude bands each day. The vertical volume mixing ratio profiles are retrieved for eight species for each event. The accumulated data for approximately 30 days cover most of the globe (limited by sunlit latitudes), and this monthly data block can be described as the volume mixing ratio of a specific species in the atmosphere as a function of latitude, longitude, and height. The data were remapped using linear interpolation for pressure levels and Gaussian weighted binning from sampling locations to a three-dimensional grid. An AVS network is constructed that allows for viewing the three-dimensional field with rendered slices at constant latitudes, longitudes or pressure levels. Discussions are given on the advantages and some disadvantages learned about from experiences applying AVS to visualize HALOE three-dimensional data.

¹Department of Earth System Science, University of California, Irvine, CA

²Office of Academic Computing, University of California, Irvine, CA

³Atmospheric Sciences Division, NASA Langley Research Center, Hampton, VA

⁴G & A Technical Software, Inc., Hampton, VA

1. INTRODUCTION

As many geophysical scientists are beginning to realize, the analysis of the fast growing atmospheric observation data set requires the support of computer software tools that allow them to view, manipulate, and display the data in a visual form before presenting a verbal description. The advantages of applying satellite-borne instruments to remotely sound atmospheric properties are that the measurements can cover a large latitude/height range and the result can be time dependent. These three-dimensional, time-dependent data sets have never been obtained before. It is important for scientists to attain access efficiently. Several software packages have been developed in recent years for scientific visualizations, and they are improving dramatically to satisfy a user's demands. In this paper, we share our experiences using one visualization package to study a three-dimensional data set.

The Halogen Occultation Experiment (HALOE), one of the nine instruments aboard the Upper Atmosphere Research Satellite (UARS) launched in September 1991, has been making measurements of a total of eight stratospheric/mesospheric chemical species at every spacecraft sunrise and sunset event (Russell et al., 1993a). We describe the HALOE observations in more detail in the next section. A HALOE monthly data block for a species can be described as its volume mixing ratio as a function of latitude, longitude, and pressure. Although HALOE does not provide daily global maps for its measured species because of the geometrical constraints described below, in most cases the monthly data block may represent the characteristic of a species for the corresponding season.

We chose the Application Visualization System (AVS) to view HALOE monthly three-dimensional data blocks. The main visualization strategy of AVS requires users to select the “building blocks” (modules) and construct a “building” (AVS network) to show the data in a way that is scientifically meaningful. The AVS provides many modules while also allowing users to write their own. This paper describes the ways we used AVS to view the three-dimensional data set. The last section discusses the advantages and disadvantages of this approach based on our experiences.

2. HALOE OBSERVATIONS AND LEVEL 3 DATA PROCESSING

By the 1993 spring American Geophysical Union (AGU) meeting, the HALOE on the UARS had already been successfully making measurements on atmospheric temperature and chemical species for more than 20 months. This experiment uses the sun as its radiation source and measures atmospheric absorptions in selected infrared spectral regions between 2.4 and 10 μm . It provides

vertical mixing ratio profiles for O_3 , CH_4 , H_2O , NO , NO_2 , HCl , HF , and aerosol extinction at every spacecraft sunrise and sunset event. More detailed information on HALOE observation principles, instrument construction, and data inversion techniques are in a paper by Russell et al. (1993a). Because UARS has a 57° inclination near-circular orbit and it takes 98 minutes to circle around the Earth at 585 km, HALOE obtains about 15 sunrise and 15 sunset measurements approximately every 24 hours. The approximately 30 sunrise/sunset tangent point (defined as the point of closest approach of a solar ray path to the Earth surface along the limb) locations for a given day are distributed in two narrow latitude bands.

Because of the complicated interplay of spacecraft path and the orientation of the Earth's limb, the time-varying sampling points are distributed in a complex pattern. Figure 1 shows HALOE daily averaged sunrise and sunset latitude positions for the 30-km tangent point altitude (TPA) as a function of time between September 8 and November 7, 1992. We select this period when the Antarctic ozone in the lower stratosphere is chemically




Figure 1. HALOE daily averaged 30-km tangent point latitude coverage for the period between September 8 (UARS day 363) and November 7, 1992 (UARS day 423).

removed from the existence of manmade chlorine compounds because it is of great interest to atmospheric scientists, as well as to the general public. During this period, HALOE observations cover from its northernmost latitude (about 75°N) to southernmost (about 78°S) in 20 to 30 days. As an example, data between September 8 and October 4 cover most of the globe, which is quite useful in studying seasonal behavior of the chemical species. Likewise, Figure 2 is the HALOE sunrise 30-km TPA positions between September 8 and October 4, plotted on the orthographic projection of the Earth's southern hemisphere. Along every latitude band for a day, 15 sunrise measurements are uniformly distributed in longitude.




Figure 2. HALOE 30-km tangent point coverage between September 8 and October 4, 1992. Sunrise measurements are available for this period in the southern hemisphere.

The HALOE data base consists of several levels of data: level 1, the raw data; level 2, retrieved temperature and the volume mixing ratios of seven chemical species and aerosol extinction as a function of pressure, latitude, and longitude at measurement locations; and level 3, gridded data that have been remapped into a three-dimensional grid of latitude, longitude, and pressure level. The data processing from level 2 to level 3 can be described as follows. First, data are interpolated onto UARS standard pressure levels (mb), defined as

$$P = 1000 \times 10^{-(i/6)}, \quad i = 0, 1, 2, \dots, 35. \quad (1)$$

The pressure spacing is $\Delta \log(P) = -1/6$, equivalent to about 2.5 km. This is about the same as HALOE level 2 data vertical spacing for its gas filter channels (CH_4 , NO , HCl , and HF). The HALOE radiometer channel's vertical spacing is about 0.3 km (O_3 , H_2O , and NO_2).

Second, on a constant pressure surface, data are averaged using a Gaussian-weighted binning method. The gridded data spacing was chosen to be 1° latitude by 2° longitude. The normalized Gaussian function applied to the data point inside a bin is

$$G = \exp \left\{ - \left[(x/\sigma_x)^2 + (y/\sigma_y)^2 \right] / 2 \right\}, \quad (2)$$

where x and y represent the differences of longitude and latitude between the data grid and the sampling position. The standard deviations are selected to be 15.0° longitude (σ_x) and 5.0° latitude (σ_y). The bin width used is two times the standard deviations in longitude and latitude, respectively. The volume mixing ratio at a grid point (Z_{ij}) is therefore calculated by

$$Z_{ij} = \frac{\sum [G \times Z] \text{ (within the bin)}}{\sum G \text{ (within the bin)}}, \quad (3)$$

where Z represents a sampling data point within the bin.

HALOE level 3 data provide three-dimensional gridded data blocks for any selected period. The use of a state-of-the-art scientific visualization package helps us view, manipulate, and analyze the data. The specific values used for parameters in the binning method can be adjusted. Several tests were performed to study the effect of these parameters on our analysis (for example, to change the values of σ_x and σ_y and bin width). Data visualization tools will no doubt add power to the tests. We will not discuss this analysis in this paper because it is not one of the main goals of the paper.

3. ABOUT AVS AND ITS APPLICATION AT THE UNIVERSITY OF CALIFORNIA, IRVINE

Advanced Visual Systems, Inc., created AVS as a software package for scientific data visualization.

Stardent Computer Corp. originally designed it for software developers designing complex software packages with graphical interfaces. The subsequent use of AVS has shown that it is quite useful as a program for scientific visualization and interpretation of data by the much larger scientific user community. At present, AVS is available on most UNIX-based computer workstations (HP, IBM, SGI, DEC, Sun, Kubota, and Data General), as well as the larger supercomputers (Convex, Cray, Thinking Machines, etc.).

AVS was originally introduced to the University of California, Irvine, through the introduction of the then-powerful Stardent Titan line of workstations. As newer versions of AVS became available on other computer platforms, the use of AVS at the university has greatly increased. Many research groups across the spectrum of research interests (medical imaging, neuroscience, combustion engineering, fluid mechanics, groundwater pollution, astrophysics, etc.), as well as geophysics, have been making substantial progress understanding their data and simulations using AVS. Several graphics laboratories have been set up using AVS on DEC, Stardent, and Convex computers.

The usefulness of AVS for scientific data visualization surfaced from the visual and data flow interfaces of AVS. The visual interface allows the system user to design the kind of visualization tool needed for his or her particular research through a series of standard point-and-click operations, which combined with the data flow structure of AVS, guide a user into building a visualization "network" from a set of simpler "module" building blocks. Using AVS for the first time therefore entails learning the visual interface functions, searching the module library for the functions that are needed, and importing one's data into an AVS format that is based on a mathematical vector field (one or more data values at points defined in a one- to three-dimensional geometry).

To most geophysical scientists, learning new computer software means investing a large amount of time. The compensation for the use of AVS is its general application across a wide variety of data. Generality allows AVS to take advantage of many types of modern visualization techniques (renderings) as well as mathematical operations (such as image processing). Most importantly, AVS allows the creation of new modules based on the

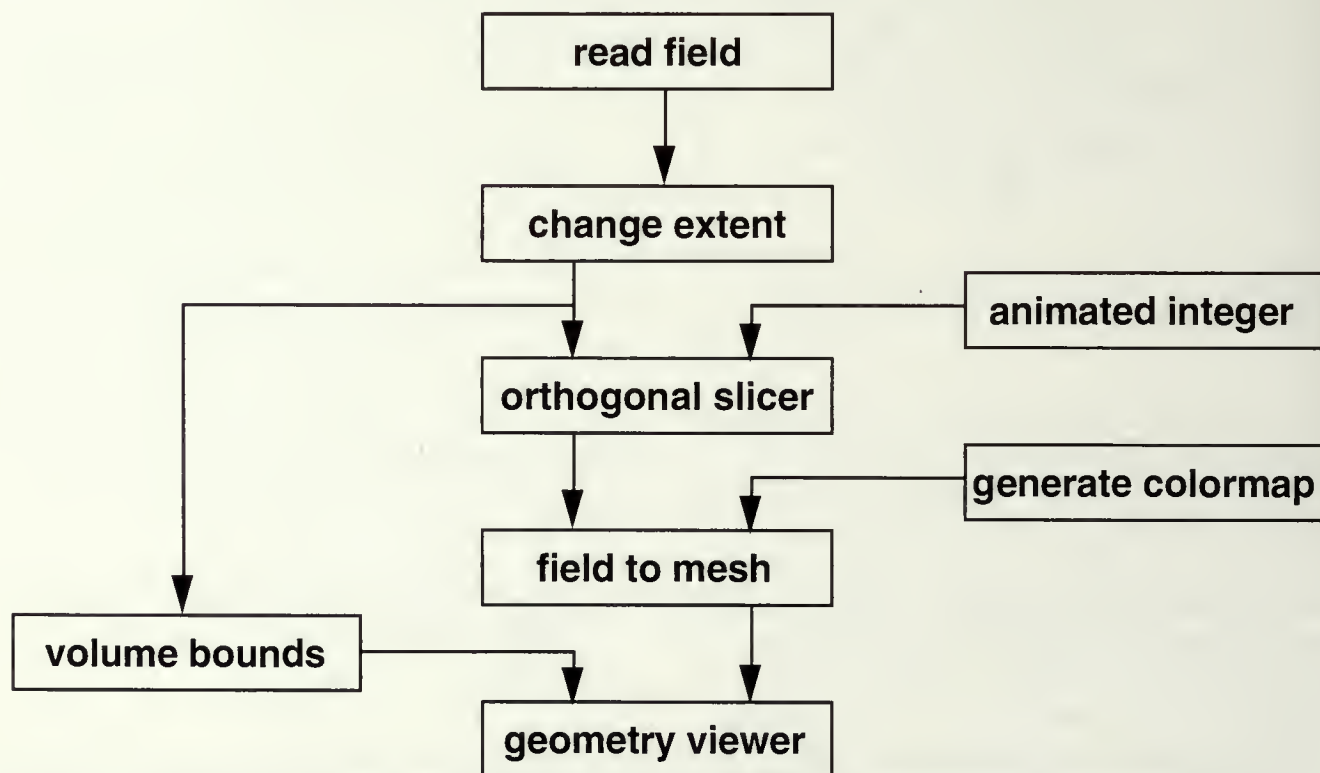


Figure 3. An example of the AVS network (connection of AVS modules) used to view HALOE three-dimensional data.

high-level computing languages C and FORTRAN, known to most scientific users.

4. THE USE OF AVS TO VIEW THREE-DIMENSIONAL HALOE DATA

As described in section 2, the HALOE data block between September 8 and October 4 covers latitudes from 72S to 72N with $1^\circ \times 2^\circ$ as latitude and longitude spacings, and we consider pressure levels between 100 (~16 km) and 1 (~48 km) mb, with $\Delta \log(P) = -1/6$ as spacing. Therefore, the total number of data points (the volume mixing ratio of ozone, for example) in this data block is $181 \times 145 \times 13$ (longitude \times latitude \times pressure).

Figure 3 shows a simple AVS network that connects basic software components (modules) needed to examine our three-dimensional data block. This network looks like a flow chart—data are read in first and then are displayed. The key module used in the network is the “orthogonal slicer,” which searches for data at selected constant longitude, latitude or pressure levels. The “generate colormap” module is used to define the color scale according to the data value range. One can turn on the “animated integer” module is used to define the color scale according to the data value range. The “animated integer” module to dynamically display the three-dimensional data, for instance, moving the orthogonal slicer at constant pressure surface through the entire range of pressure levels with the desired number of steps. The explanation of other modules in Figure 3 can be found in the AVS user’s guide, the AVS on-screen help menu, or the following examples.

HALOE level 3 CH_4 data are examined using the AVS multi-orthogonal slice module, as shown in Figure 4. Figure 4(a) illustrates three ways to examine the data block, placing a rendered slice at any desired constant latitude, longitude, or pressure level. Figure 4(b) is more interesting scientifically because it shows two rendered slices of CH_4 mixing ratios at southern and northern high latitudes, respectively. CH_4 is often treated as a dynamical tracer for transport processes in the Earth’s atmosphere (more specifically, stratosphere extending from 100 mb to 1 mb) because it is not photochemically produced in the atmosphere and its stratospheric lifetime is about the same order of magni-

tude as the time constants for transport by the meridional winds. At northern high latitudes, CH_4 is uniformly distributed along a latitude circle. In contrast, CH_4 shows strong longitudinal asymmetry at southern high latitudes in austral early spring, the period that we selected. It is important for atmospheric scientists to monitor the behaviors of “tracer gases” such as CH_4 during the course of Antarctic winter and spring. The evidence of low CH_4 within the polar vortex in the lower stratosphere, where ozone chemical destruction occurs, indicates that the air inside the vortex experienced significant descent prior to the observation.


To properly orient the data to geographical features, we created an AVS module (using FORTRAN) to convert the input data to an AVS field where each data value was mapped to the corresponding x,y,z position based on its latitude, longitude, and pressure level [the radial spacing of the data value is set by an adjustable scale factor multiplying the pressure index (i in equation 1)]. While the resulting pressure level surface is merely a relative measure of the real pressure level, the latitude/longitude information is directly comparable to the continental geometry. Among the many types of possible renderings of this remapped data, we mainly used constant pressure surfaces (see Figure 5) because their comparison with land masses is obvious.

A useful AVS function is general object manipulations. In our case, the global ozone data at a constant pressure level can be manipulated simply by clicking and moving the computer mouse. Figure 5 shows two snapshots of our “object” while we moved it around. With a little practice, users can easily learn three ways of manipulating the object—translation, rotation, and scaling. Through rotating our ozone surface, we could examine global ozone data and make comparisons between different latitude regions.

HALOE data show that the ozone-depleted area is strongly correlated with southern polar vortex location. Here, the “tracers” measured by HALOE show air descending, and this dynamically isolated and chemically disturbed area is not zonally symmetric over the pole in late September and early October 1992. The “ozone hole” has a large geographic coverage extending to the southern tip of South America. In the Northern Hemisphere,

HALOE CH₄ MIXING RATIO SEPT.8 – OCT.4, 1992

(a)



(b)




Figure 4. HALOE-measured CH₄ volume mixing ratios in the period between September 8 and October 4, 1992—(a) an example of using rendered slices to view the data block at arbitrary constant latitude, longitude, or pressure level; (b) two slices of CH₄ at 71°S and 71°N, respectively, showing the different characteristics of CH₄ along the two latitude circles.




Figure 5. HALOE-measured ozone volume mixing ratios at 32 MB in the time period between September 8 and October 4, 1992. AVS allows users to manipulate the "object" (rotation and translation, etc.) by clicking and moving the computer mouse. Two snap shots of global ozone data are shown here. In the top image, the outline of South America is just above center. The bottom image looks down on North America.

however, HALOE ozone data show zonal symmetry when summer turns to winter. More detailed scientific analysis of HALOE measurements in polar regions appears in other publications (such as, Russell et al., 1993b; Tuck et al., 1993).

5. DISCUSSION

With some imagination, AVS provides many useful tools to visualize three-dimensional data. Before the 1993 AGU spring meeting (May 26—28) when this book was initiated, we used AVS version 4.1 on a DEC 5000/200 workstation with 32 MB memory. Although AVS is also available on the Convex machine at the University of California, Irvine, we found we had to share that resource with many other users and the processing time was quite slow. AVS 5.0 was released before the AGU meeting, and it recently became available. This new version of AVS has several major improvements. For instance, a color legend module is added to the standard AVS module library. This module is useful for studying data quantitatively.

AVS is also useful for viewing and manipulating large three-dimensional satellite data sets. The outputs from mass production of satellite data are data files containing valuable information for scientific studies. For example, since launch in September 1991, UARS, with nine scientific instruments onboard, has been making atmospheric and space environment measurements over two Antarctic winter/spring (1991 and 1992) and two Arctic winter/spring (1992 and 1993) periods, and it has already started its 1993 Antarctic winter/spring measurements. The concentrations of many atmospheric species have never been observed prior to UARS. An efficient way to view those data is badly needed, which not only allows scientists to scan over a large amount of data in a relatively short period, but also facilitates a quick publication of the analysis. As we showed in previous sections, the two AVS networks we constructed, combined with on-screen object manipulating functions of AVS, allow us to examine the data efficiently.

With a little help, atmospheric scientists can learn a few basic functions (software components) that are commonly used in their field. AVS provides hundreds of standard modules, but only a few are useful for a specific case in a scientific sub-field to display the data

in a meaningful way. AVS also provides user-friendly control panels to allow users to change adjustable parameters in a module, as well as viewer control panels to allow users to manipulate objects and to optimize the viewing effect. Learning a new visualization package and getting familiar with it are always time consuming and work intensive, but it is quite enjoyable to work on AVS with the HALOE data.

Although AVS allows users to write their own modules, it is not usually the preferred choice. While AVS modules are actually C functions or FORTRAN subroutines, they contain enough AVS-based function calls to make the subroutine creation process intimidating enough for the casual programmer. However, when needed, the ability to build a module can be easily expanded to create modules that can manipulate the scientific data, imagery of that data (without resorting to X window programming), or even three-dimensional geometries from that data (without any knowledge of the underlying graphics language on that platform). A substantial number of such user-generated modules are stored at the International AVS Center at the North Carolina Supercomputer Center. Such modules can be obtained from the center via anonymous ftp to [avs.ncsc.org](ftp://avs.ncsc.org).

There were a few disadvantages. A new software package will always have initial defects, which are eventually overcome in later releases. There are some problems to be considered before investing time and effort.

The major problem was the processing rate for our data, which was $181 \times 145 \times 13$ (longitude \times latitude \times pressure), and was very slow using a DEC 5000/200 workstation (23 spec mark) with 32-MB memory. The graphic board that the DEC station uses is a PXG board, which has a speed of approximately 55,000 shaded polygons/second. The rotation of our “object” (see Figure 5) is slow but tolerable. The data rendering, for instance, changing from one constant pressure surface to another, is quite slow (typically a couple of minutes). We assume that the 32-MB memory is not adequate (swapping between the CPU and its hard disk is usually the problem). A faster CPU speed and larger memory (larger than 32 MB) is recommended for similar size data sets.

AVS is a very powerful tool for viewing three-dimensional data, but from our experiences, it is not

a good package for two-dimensional or one-dimensional plots (current version). The labeling ability of AVS is relatively poor, which not only affects the quality and readability of two-dimensional plots but also causes difficulty in quantitative study using three-dimensional display. Another constraint is that users have to use AVS line commands to group their selected objects and manipulate the object group (rotation/translation). Although writing one's own module could solve many or all of the above problems, many individual users will lose interest in using this package for some of their studies. Later AVS versions could accommodate more user demands with standard functions. The linkage of AVS with other mathematical and visualization tools (e.g., IMSL and IDL) obviously adds power to this multi-software package, although we have not used these combinations on our data set.

Acknowledgments. The authors would like to express thanks to the HALOE science and flight operation teams for providing an excellent satellite data set and for encouraging the exploration of a

variety of visualization tools. We would also like to thank colleagues at the Office of Academic Computing, University of California, Irvine, for their enthusiastic support, from hardware configurations to software lessons and consultations. This work is sponsored by NASA grant NAS1-19155.

REFERENCES

- Russell III, J.M., L.L. Gordley, J.H. Park, S.R. Drayson, W.D. Hesketh, R.J. Cicerone, A.F. Tuck, J.E. Frederick, J.E. Harries, and P.J. Crutzen, The Halogen Occultation Experiment, *J. Geophys. Res.*, "Special Issue on UARS," 98, 10,777-10,797, 1993a.
- Russell III, J.M., A.F. Tuck, L.L. Gordley, J.H. Park, S.R. Drayson, J.E. Harries, R.J. Cicerone, and P.J. Crutzen, HALOE Antarctic Observations in the Spring of 1991, *Geophys. Res. Lett.*, 20, 719-722, 1993b.
- Tuck, A.F., J.M. Russell III, and J.E. Harries, Stratospheric Dryness: Antiphased Desiccation Between Micronesia and Antarctica, *Geophys. Res. Lett.*, "Special Issue on UARS," 20, 1227-1230, 1993.

A Versatile System for Processing Geostationary Satellite Data With Run-time Visualization Capability

M. LANDSFELD, C. GAUTIER and T. FIGEL

Earth Space Research Group, University of California, Santa Barbara, CA

To better predict global climate change, scientists are developing climate models that require interdisciplinary and collaborative efforts in their building. We are currently involved in several such projects but will briefly discuss activities in support of two such complementary projects: the Atmospheric Radiation Measurement (ARM) program of the Department of Energy and Sequoia 2000, a joint venture of the University of California, the private sector, and government agencies. Our contribution to the ARM program is to investigate the role of clouds on the top of the atmosphere and on surface radiance fields through the data analysis of surface and satellite observations and complex modeling of the interaction of radiation with clouds. One of our first ARM research activities involves the computation of the broadband shortwave surface irradiance from satellite observations. Geostationary satellite images centered over the first ARM observation site are received hourly over the Internet network and processed in real time to compute hourly and daily composite shortwave irradiance fields. The images and the results are transferred via a high-speed network to the Sequoia 2000 storage facility in Berkeley, where they are archived. These satellite-derived results are compared with the surface observations to evaluate the accuracy of the satellite estimate and the spatial representation of the surface observations. In developing the software involved in calculating the surface shortwave irradiance, we have produced an environment whereby we can easily modify and monitor the data processing as required. Through the principles of modular programming, we have developed software that is easily modified as new algorithms for computation are developed or input data availability changes. In addition, the software was designed so that it could be run from an interactive, icon-driven, graphical interface, TCL-TK, developed by Sequoia 2000 participants. In this way, the data flow can be interactively assessed and altered as needed. In this environment, the intermediate data processing "images" can be viewed, enabling the investigator to easily monitor the various data processing steps as they progress. Additionally, this environment allows the rapid testing of new processing modules and allows their effects to be visually compared with previous results.

1. INTRODUCTION

To aid in the prediction of global climate change, scientists are developing climate models that require interdisciplinary and collaborative efforts in global data acquisition, complex model building, and innovative result analysis and interpretation. Scientific teams, distributed worldwide, analyze and produce fields of geophysical parameters used as input to climate models. These parameters include sea surface temperature, surface albedo, carbon flux, wind fields, atmospheric temperature, and moisture profiles, among others. To accomplish this enormous task, data need to be easily transferred between sites via high-speed computer networks at regular intervals, processed at various sites, and finally transferred to data base storage facilities for other scientists to utilize. We have developed a system for producing global fields from geostationary satellite data and visually assessing the results as a quality control measure. This system computes surface solar radiation flux, but the approach can be generalized and applied to any computer model, especially satellite data processing systems.

2. DATA SETS AND INFRASTRUCTURE

We have integrated the capabilities of two distributed projects in which we are involved to produce the computation and display system described herein: the ARM (DOE, 1990) program of the Department of Energy and Sequoia 2000 (Dozier, 1993), a joint venture of the University of California, the private sector, and government agencies. These are both long-term projects; we have developed the computer code of our radiative transfer modeling system with this in mind.

ARM is a multi-investigator research endeavor aimed at studying the role of clouds in climate through the improvement of their parameterization in numerical models. A 10-year field program is being developed to collect a comprehensive data set to characterize cloud properties and their interaction with the environment. The data set is collected from a number of high data accumulation rate instruments, such as surface interferometers, satellite imagers, and atmospheric sounders.

Our contribution to the ARM program is to investigate the role of clouds on the top of the atmosphere and on surface radiation fluxes through the data analysis of surface and satellite observations and the complex modeling of the interactions between solar radiation and clouds. One of our first ARM research activities involves the computation of a shortwave (0.25 - 4.0 μm) solar radiation flux at the surface from satellite observations, and we use this task as the example for our visualization system. Geostationary satellite images from GOES, centered over the first ARM Southern Great Plains (ARM-SGP) observation site (located in Oklahoma), are received hourly over the Internet network and processed in real time to compute hourly and daily surface shortwave irradiance. Every hour, images are received and rectified to ground coordinates. These images and the results of our radiative transfer model (RTM) are transferred via a high-speed network to the Sequoia 2000 storage facility in Berkeley, where they are archived and made available to other scientists. In parallel, we receive in-situ surface shortwave flux measurements from the ARM-SGP site pyranometer for model validation.

Sequoia 2000 is an interdisciplinary program bringing together global change and computer scientists to address issues of rapid data transfer, data management, and advanced visualization techniques in support of global change research. Global change science, in turn, serves as a test bed for the advanced distributed information systems developed by the computer scientists. The system provides the infrastructure to enable scientific interactions among researchers of different disciplines and among researchers employing different methodologies. The program has developed an information system, not just a data system. For example, the Sequoia 2000 data base provides

geophysical and biological information, not just raw data from space-borne instruments or in-situ sensors.

3. RADIATIVE TRANSFER MODEL AND ALGORITHM

At the heart of the visualization system is the RTM, which produces the data layers for the graphical user interface (GUI). The simple model used was first introduced by Gautier et al., (1980), with later improvements from Diak and Gautier (1982). The following is only a general description of the model because the purpose of this paper is to describe our data handling, algorithm design, and visualization techniques. A flowchart for the current processing algorithm is displayed in Figure 1.

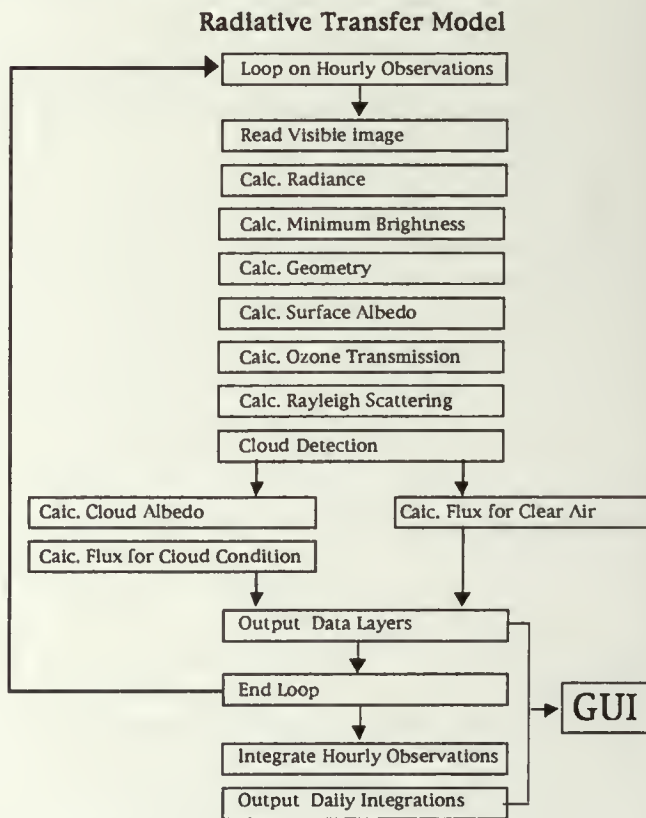


Figure 1. Flowchart of the radiative transfer model.

Daily average downwelling surface solar radiation fluxes are produced from visible geostationary satellite images and are recorded hourly for our experiment. As the first step, the processing system ingests the visible satellite data and calculates both solar and satellite geometries

(zenith and azimuth angles). Upwelling radiance values are then computed using satellite calibration coefficients obtained from Whitlock (personal communication, 1994). A surface albedo is estimated by computing the minimum brightness values from a set of images taken at the same hour each day over the previous 15 days. This “minimum brightness” compositing technique is quite common in remote sensing. The minimum brightness value is assumed to correspond to a “clear sky” value for each pixel. The directional surface reflectance is calculated from this value and is used to estimate the surface broadband albedo using the spectral and directional transformation characteristics of the surface type over which the solar flux is estimated.

Two atmospheric parameters are calculated next: ozone transmittance and Rayleigh scattering. Subsequently, a cloud detection threshold value is computed, based on the clear sky pixel values, to distinguish between clear and cloudy pixels. A computation for downwelling surface solar flux is made for either the clear or cloudy conditions, based on the cloud detection results. If the condition is determined to be clear, a clear sky radiative transfer model is applied, which computes the attenuation of the solar flux by water vapor and ozone absorption and by Rayleigh and aerosol scattering using climatological concentrations of the different absorbers and scatterers. If the condition is determined to be cloudy, cloud albedo directional reflectance is computed and used as the cloud broadband albedo. The surface solar flux is calculated using the cloud albedo to estimate cloud transmission and absorption and then applying this to the computed downwelling clear flux. After the hourly observations are processed for each pixel composing the image, the hourly estimates are integrated over time to obtain daily averaged estimates.

The satellite-derived results are compared with surface observations to evaluate the accuracy of the satellite-based estimates. Comparisons are performed on a 3-km-by-3-km spatially averaged model estimate to compensate for locational inaccuracies of the satellite observations and spatial integration (2 solid angle) of the surface observations. To further adjust for these effects, the surface observations are time averaged over 20-




Figure 2. In-situ shortwave measurements versus radiative transfer model estimations.

minute intervals encompassing the satellite observation time—the average time for clouds to move across the 3-km field of view. Averages from approximately 800 observations are compared with the RTM estimates and are displayed in Figure 2. The model produces very good results, with a mean difference of about 45 W m^{-2} and a standard deviation of the mean difference of 53 W m^{-2} , or about 10 percent of the mean value. The mean difference indicates that there is a small bias in the satellite estimates, which might result from either the satellite sensor calibration, the surface pyranometer calibration, or the model assumptions. With high-quality surface instrument calibration, these comparisons could be used to improve the satellite sensor calibration.

4. SOFTWARE SYSTEM

As mentioned previously, this RTM is a long-term project and in the future will be continually refined, modified, and improved. To facilitate this, we have developed, through the principles of modular programming, a software system that is easily modifiable as new algorithms for computation are developed or as new observational inputs become available. The algorithm functions have been developed to be coherent, self-contained modules that perform calculations for individual

radiative transfer properties, such as Rayleigh scattering or surface reflectance. To maintain a simple, logical program flow, all principal algorithm functions are called from the main program and all “accounting” procedures, such as time tracking, are executed in the main program. This leaves the modules free of extraneous activities and allows them to be easily modified or replaced with new modules in the future. Furthermore, all the parameters are passed to each function through the argument list, thereby preventing the use of global variables that unnecessarily interconnect the modules. As model development progresses, the modules can be “plugged” in and out without the burden of algorithm accounting details affecting the modules’ performance. As an example, the cloud albedo routine does not require the current processing observation time to calculate solar geometries, but instead is passed the geometric quantities through its argument list and requires only the number of observation pixels to process.

All the above-mentioned routines produce a “data layer,” or intermediate data set, that is used by other routines—that is, each function produces a parameter for the entire hourly image and passes that computation to other functions. This “data layer” approach allows for a clean style of programming, but is memory intensive because each layer contains the entire field that is being processed in the run. Efforts are currently under way to enable the program to process smaller segments of the images with a user-definable segment size. Again, all these accounting procedures will be independent of the modules and executed in the main program.

Written in the C programming language, the RTM algorithm has been developed with command line options, and the *getopt* command is used to parse the options on the command line. We have developed the RTM algorithm so that all the intermediate data sets can be output for visualization with a command line “flag” selection. If the option flag for a data layer is chosen, such as “-a” for surface albedo, the layer is written to the computer disk and can subsequently be displayed or analyzed by other routines. For example, the user can set such options as day of year to process and select which type of output to generate with the command line statement

RTM -Y 93 -D 36 -s -a

setting the day to process to year 1993 (-Y 93), day of year to 36 (-D 36) and requesting the outputs, shortwave (-s), and surface albedo (-a). This design feature allows for a convenient coupling with the display routines.

4.1 Display Routines

Our data visualization is performed, in part, with Interactive Data Language (*IDL*), a graphical display software [Research Systems, Inc., 1991]. It is a commercially available interactive software system and programming language designed for plotting and visualizing images of scientific data. The language contains simple syntax routines for displaying data that can be combined to produce highly customized display programs for the visualization of the RTM products. *IDL* provides multiple plotting capabilities so that several images can be displayed simultaneously in the same window as in Figure 3. This display configuration allows a side-by-side comparison of the parameters that are produced in the processing of the model. In this example, six data layers are displayed in the panel. In Figure 3(a), the model’s input, GOES-7 brightness field is displayed. Also displayed with each data layer is a self-scaling legend and the average value of that parameter as computed by the *IDL* display routine. Figures 3(b-e) display some of the intermediate processing steps generated by the RTM—namely, the solar zenith angle, cloud detection results, and the estimated cloud and surface albedos, respectively. The resulting downwelling shortwave flux to the surface is displayed in Figure 3(f). With this graphical configuration, we can visually inspect the step-by-step progress of the model. By examining the value range in the legend, or the statistics, we can critique any changes in the results as compared to previous runs. These composite visualization images can then be read in as frames for an animation sequence and played back at varying speeds on the computer screen or stepped through on a frame-by-frame basis with the *IDL* animation tool. In this manner, the model’s step-by-step progression can be examined for each processing day.

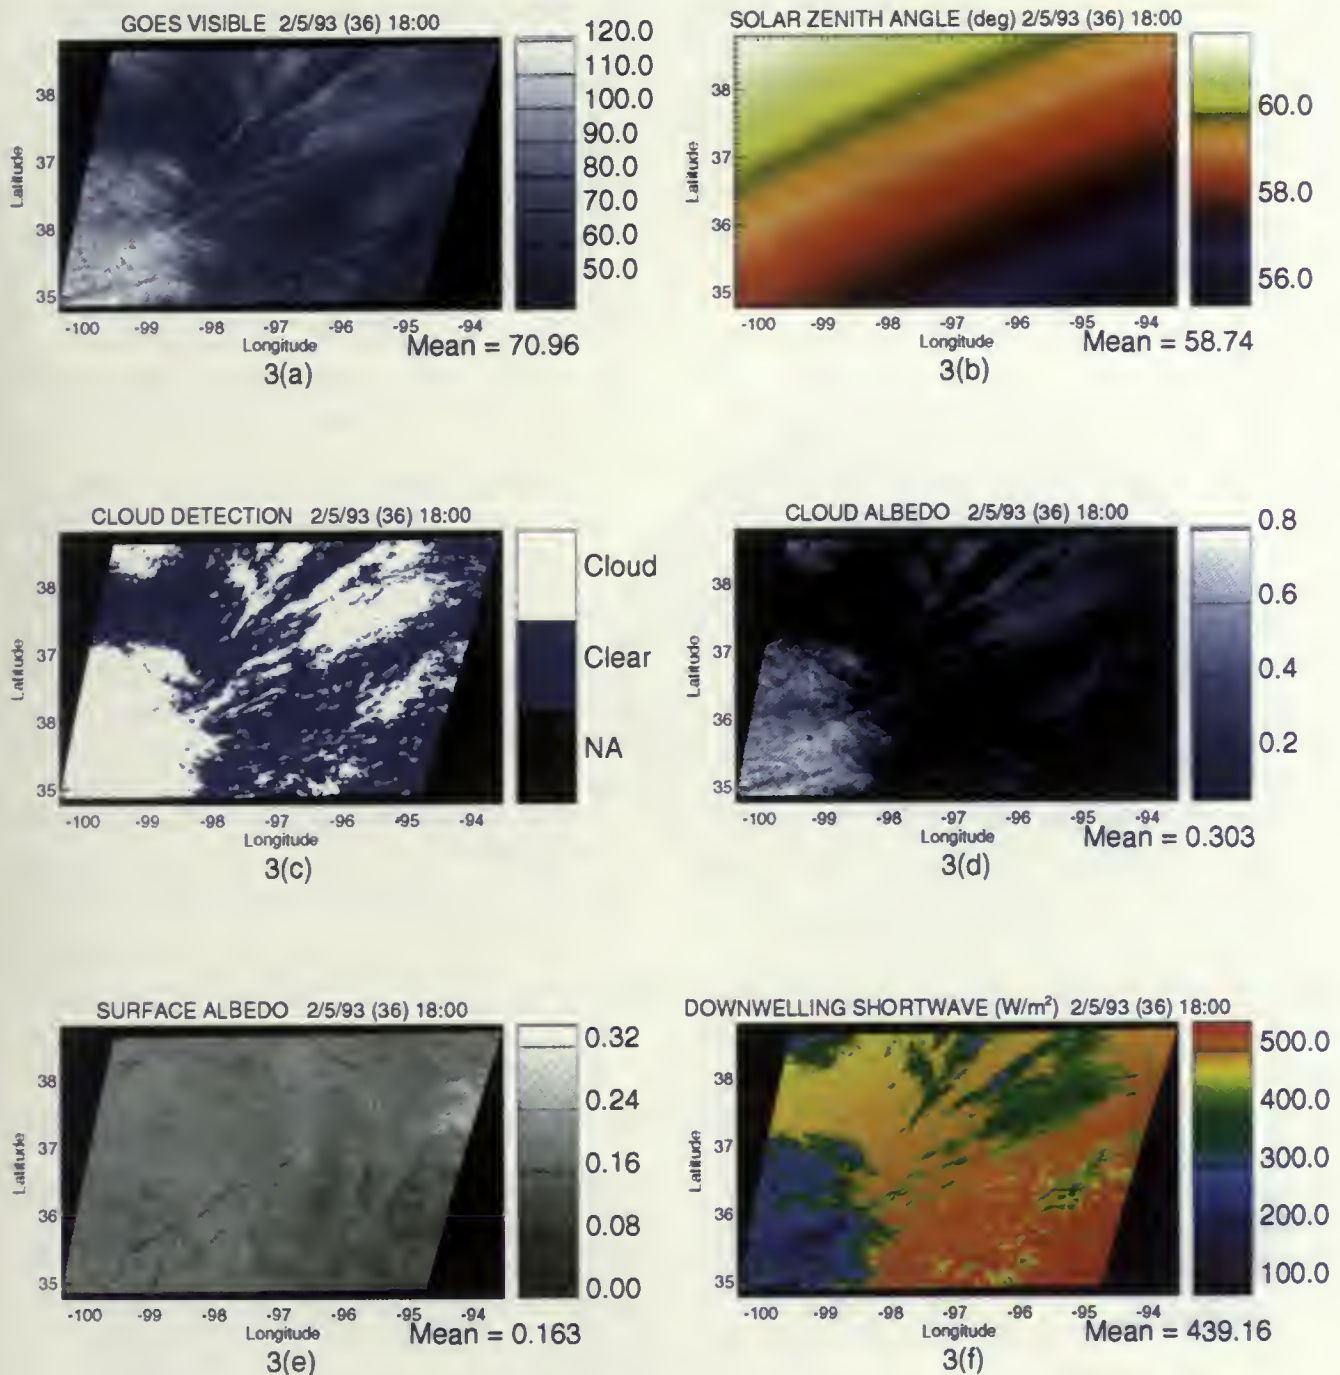


Figure 3. Intermediate data sets displayed by the visualization system.

4.2 Graphical User Interface

TCL-TK is the GUI used in our software system. It was developed by Sequoia 2000 participants (Ousterhout, 1993). TCL-TK is a shell language GUI builder, and the syntax is similar to the UNIX shell programming language. It is public domain software available via anonymous ftp at sprite.cs.berkeley.edu. It runs on various UNIX

platforms (Sun, DEC, IBM, HP, and SGI), as well as PC platforms (SCO, UNIX, and LINUX). The source code uses about 23 Mb of disk space to compile, and it runs on a PC with 8 Mb of RAM.

We found that programming in TCL-TK was preferable to developing GUI's using Motif or X11 library functions in C programs. The programmer can write scripts that are about 10 percent of the size of the C programs, with a similar order of

magnitude in CPU execution time savings. Public domain scripts written in TCL-TK can be used to build TCL-TK applications that use a mouse-driven selection mechanism. This capability allows users who do not know the command set to create their own interfaces without the overhead of learning the script language.

An example of the TCL-TK interface is displayed in Figure 4. The interface has a button orientation similar to many menu-driven software packages available today. Buttons can be programmed to execute processes or execute display routines, open other windows, toggle options, or select values for options. The TCL-TK interface to the RTM algorithm allows the user to visualize model results as it is executing. This is accomplished by setting up two named pipes (or processes). The first pipe is connected to the RTM algorithm and, as the algorithm is executing, it writes images to disk files as they are processed during each iteration. The TCL-TK interface oversees this process. When the files are produced, the necessary commands are written to the second pipe, an *IDL* process, which then reads and displays the images as they are produced.

The TCL-TK interface sets up the command line statement for execution, as described above, and performs the execution upon selection of the process button. The TCL-TK interface provides an easy way for users to select the parameters to use as input to the model. A user can use widgets to select various options and manipulate slide bars to choose values for those options. Examples of all of these modes are displayed in Figure 4. In the upper left corner is the execution button *Execute RTM*. Selection of this widget, with the mouse, starts the execution of the RTM after all other options are set. Directly below this button, the *Processing Options* button opens the *Processing Output Options* win-

dow displayed in the lower right corner of Figure 4. In this window, up to six of the listed parameters can be selected to be written out and hourly products displayed during processing (as in Figure 3). When the options are selected, the checkboxes immediately to the left of the names are highlighted. The *Processing Machine* options allow the user to select the computer with which to execute the RTM. In the upper left window panel is a button for *Demo Mode*, in which the model is not executed but instead only the *IDL* display routines are performed. The *Demo Mode* is a method for easy review of previously processed data layers.

Below this are three slider selectors; the first two produce inputs for the command line options to the RTM, and the third produces input for the five "*Display ...*" buttons below it. The two images in Figure 4 are examples of these "*Display ...*" buttons. The lower left image is the daily integration of downwelling shortwave over the ARM-SGP site, and the upper left image displays the results of another RTM to a global data set. The *Display Lat, Lon and Value* button executes an *IDL* command that opens another window and displays the associated latitude, longitude, and image values for the location in the image over which the cursor is located.

Table 1 shows an example of the TCL-TK code used to produce the *Execute RTM* button, which illustrates its similarity to UNIX shell programming.

4.3 Visualization System

In developing the software involved in calculating the surface solar radiation flux, we have produced an environment whereby we can easily monitor and modify the data processing as required. Occasionally, the images that are received over the network contain missing or corrupted data. With our visualization system,

Table 1.

<pre> button .RTM -text "Execute RTM" -command RTM proc RTM {mode year doy options} { if { (\$mode=="Demo") } { set options "\$opti ons-D" } set \$Sh_pipe [open {/bin/sh } w+] set command " \ cd \$process_dir\; RTM -Y \$year -D \$doy \$options\ " puts \$Sh_pipe \$command Display_hourlies } </pre>	<pre> # define RTM button # procedure to run the RTM # set option value for Demo mode # open pipe to shell # set value of command string # print command to shell pipe # execute display routine </pre>
--	---

these defective images can be visually detected and corrected, if feasible, or eliminated from the processing stream. As the model produces results for the intermediate data steps, they are displayed in a window (as in Figure 3) until the next iteration of the model is completed. The image then is read in as a frame for the animation routine, and the next

iteration's set of selected images is displayed. When the model finishes executing, the images that have been saved in the buffer are sequenced together automatically for review by the user. We use this system to play back the series of images that the RTM algorithm produces over the course of its execution. The animation is useful for examining the images in a

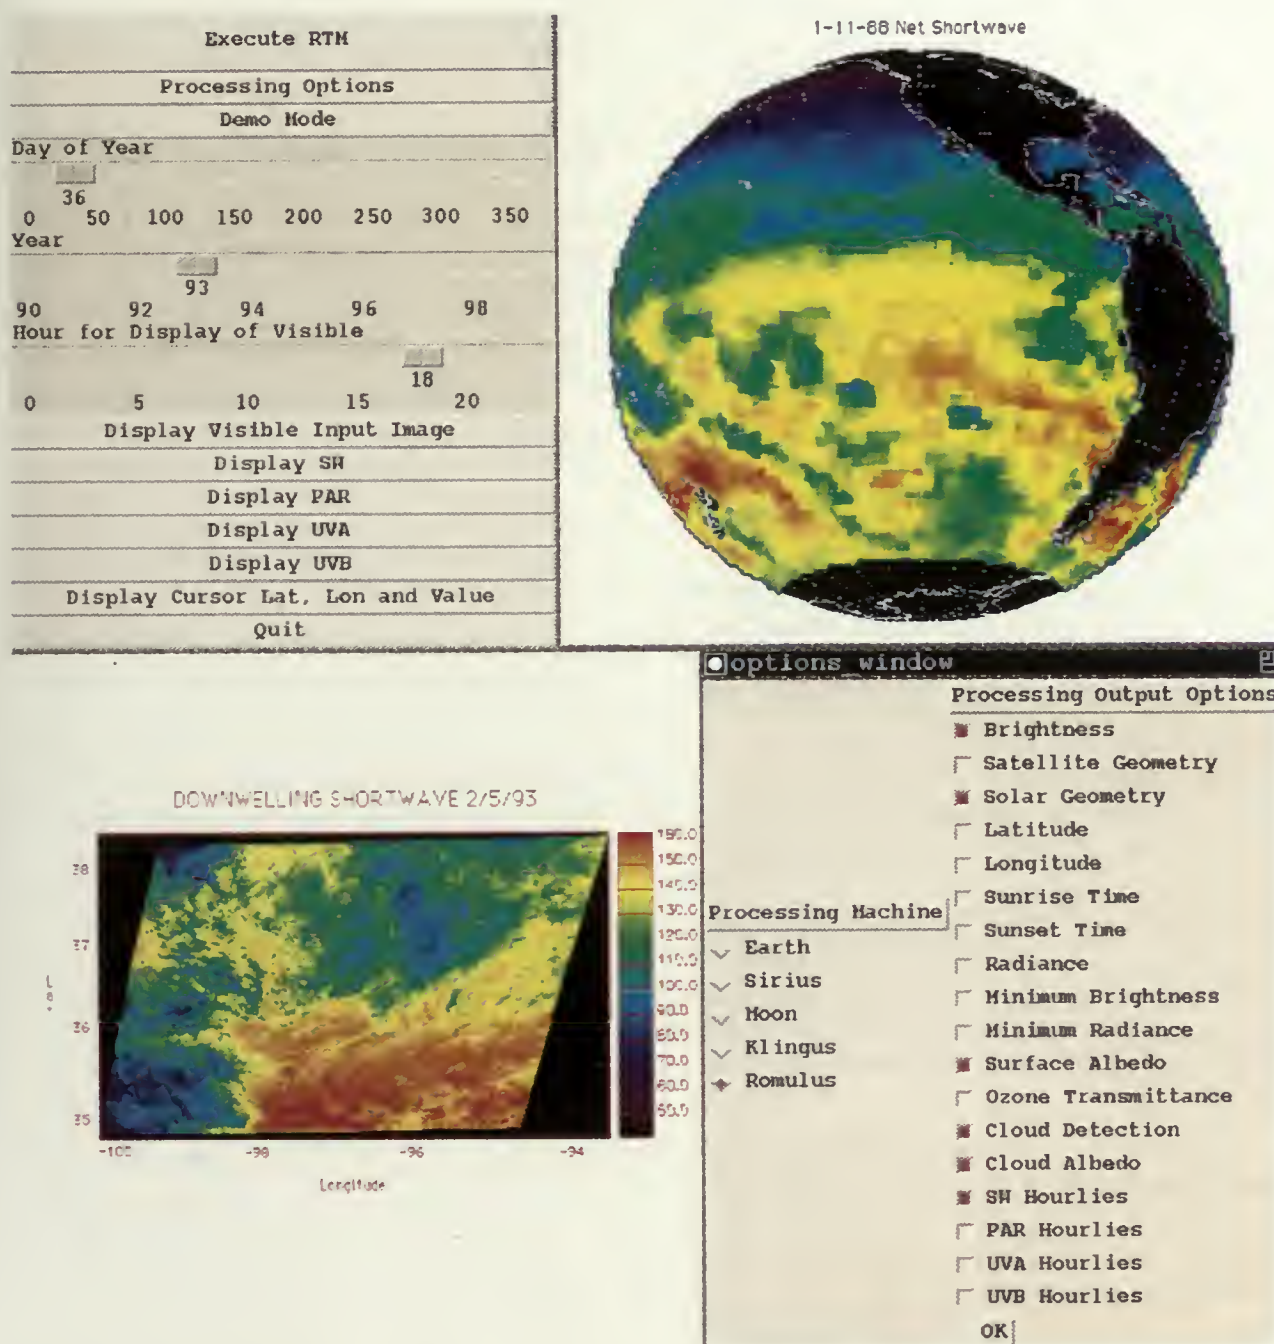


Figure 4. Visualization system TCL-TK interface (upper left and lower right), GOES visible channel over the ARM-SGP site (lower left), and a global application of another radioactive transfer model.

time sequence that can reveal error patterns not usually seen in static image evaluation. For example, upon examining some of the first outputs of the RTM algorithm, we observed dark patches in the surface albedo data layers. When these were sequenced together, a pattern developed that was recognized as cloud shadows moving across the images. We have since developed a routine for removing these effects from the surface albedo map.

The TCL-TK interface can be programmed to open other windows with multiple options. This promotes the convenient organization of display routines that are developed in the *IDL* language and used to display various inputs or outputs during model development. Menus are opened, and a number of display routines or other processing programs and their associated option selection widgets can be organized for easy execution. This is effective during the development stage of the model and the daily processing of the data streams received over the network.

5. CONCLUSIONS AND FUTURE PLANS

The TCL-TK programming language was easy to use for developing complex packaging of display routines. This capability allowed us to design a convenient method for displaying our RTM algorithm results during execution and provided a means for interactively monitoring our RTM algorithm as it is developed. The results reported here are for limited size images (400 x 700 pixels), but we are in the process of developing a similar system for global images produced by the International Satellite Cloud Climatology Project (ISCCP) [Schiffer and Rossow, 1983]. This system will compute the surface solar flux, photosynthetically available radiation (PAR), and ultraviolet A and B for global images.

Acknowledgments. The software development presented here has been accomplished with funding from the Department of Energy through the ARM program (DOE grant # DE FG03-90 ER61062) and with computer equipment provided by the Digital Equipment Corporation as part of the Sequoia 2000 project.

REFERENCES

- Diak, G.R. and C. Gautier, Improvements to a Simple Physical Model for Estimating Insolation from GOES Data, *J. Climate and Appl. Meteo.*, 22, 505-508, 1982.
- DOE, *Atmospheric Radiation Measurement Program Plan*, U.S. Department of Energy, Office of Energy Research, Office of Health and Environmental Research, Atmospheric and Climate Research Division, DOE/ER-0441, available from NTIS, Department of Commerce, Springfield, VA 22161, 1990.
- Dozier, J., How Sequoia 2000 Addresses Issues in Data and Information Systems for Global Change, in *Energy and Water Cycles in the Climate System*, edited by E. Raschke and D. Jacob, NATO ASI series, Series I, Global Environmental Change, vol. 5. Springer-Verlag, Berlin, 1993.
- Gautier, C., G.R. Diak, and S. Masse, A Simple Physical Model to Estimate Incident Solar Radiation at the Surface from GOES Satellite Data, *J. of Meteo.*, 19, 1005-1012, 1980.
- Ousterhout, J.K., *An Introduction to TCL and TK*, Addison-Wesley Publishing Company, Inc., Draft Version, October 1993.
- Research Systems, Inc., *Interactive Data Language, Users Guide*, Research Systems, Inc., Boulder, CO, 1991.
- Schiffer, R.A. and W.B. Rossow, The International Satellite Cloud Climatology Project ISCCP: The First Project of the World Climate Research Program, *Bull. Amer. Meteor. Soc.*, 64, 779-784, 1983.

Visualization of Stratospheric Ozone Depletion and the Polar Vortex

LLOYD A. TREINISH

IBM Thomas J. Watson Research Center, Yorktown Heights, NY

Direct analysis of spacecraft observations of stratospheric ozone yields information about the morphology of annual austral depletion. Visual correlation of ozone with other atmospheric data illustrates the diurnal dynamics of the polar vortex and contributions from the upper troposphere, including the formation and breakup of the depletion region each spring. These data require care in their presentation to minimize the introduction of visualization artifacts that are erroneously interpreted as data features. Nongeographically registered data of differing mesh structures can be visually correlated via cartographic warping of base geometries without interpolation. Because this approach is independent of the realization technique, it provides a framework for experimenting with many visualization strategies. This methodology preserves the fidelity of the original data sets in a coordinate system suitable for three-dimensional, dynamic examination of atmospheric phenomena.

1. INTRODUCTION

In the Earth and space sciences, it is very common to organize geographically located data as a rectilinear grid with horizontal extent over the entire surface of the Earth (i.e., latitude and longitude). In two dimensions, this implies a topological primitive that is a rectangle of various sizes. In three dimensions, the cell is a parallelepiped, with the height corresponding to altitude or atmospheric pressure, for example. These rectilinear mesh structures are ill-suited for the study of phenomena that occur continuously over a nominally spherical surface (i.e., they tear the data). In addition, these grids may not be fully populated because of missing data (such as when observations could not be made). In such cases, the grids could be viewed as being irregular.

Independent of grid type, cartographic techniques are often introduced to suitably deform the data to compensate for the problems inherent in the original structure (i.e., the use of a rectilinear representation for a spherical surface). Traditionally, such a transformation is accomplished by defining a new Cartesian grid in the cartographic projection coordinate system and then interpolating from the original rectilinear grid to the new one prior to any other operation. Given the curvilinear nature of such transformations, nonlinear interpolation techniques are typically required to make the transformation of acceptable quality. Figure 1 illustrates this operation schematically, in which the data values at three points in the original rectilinear, regular grid are combined to define the value at a single point in transformed space (e.g., the values are weighted averaged, where the weight is a function of distance from their original points to that in transformed space). In addition to being




Figure 1. Data transformation by interpolation.

computationally expensive, such interpolation makes it impossible to preserve the fidelity of the data prior to rendering, especially if regions of no data or other discontinuities are present. Such discontinuities would be smoothed out.

Alternatively, by warping the underlying mesh structure, the geometry itself is transformed without affecting the data. Figure 2 illustrates this operation schematically in which there is a one-to-one mapping of each point in the original and the transformed grid. At each node in the deformed grid, there is a data value that corresponds to a specific node in the regular grid.




Figure 2. Grid transformation by coordinate warping.

Thus, any realization (the application of one or more visualization strategies that generate renderable geometry from a collection of data) is independent of the choice of a specific cartographic coordinate system, the data or mesh themselves, or how the data are specified with respect to the underlying mesh (e.g., assigned at each node or to an entire cell at its center). In addition, interpolation is not required as the initial operation to be applied to the data to be visually correlated. Instead, interpolation can be isolated to be the last step in the visualization process, namely rendering (e.g., Gouraud-shaded surfaces).

The use of appropriately warped curvilinear

grids can preserve the fidelity of the data prior to rendering. This requires an environment that supports direct realization and rendering of data on curvilinear grids as well as regular ones. This must be coupled with the ability to independently manipulate data and its underlying mesh structure or base geometry. In addition, the ability to simultaneously render disparate geometry (such as points, lines, surfaces, and volumes of varying color and opacity) is helpful in viewing the realization of multiple data sets.

Correlative visualization implies two tenets. The first is the capability to look at multiple sets of data in exactly the same fashion (visual comparison within a common framework)—that is, displaying different sets of data of disparate structure in an arbitrary geographic coordinate system independent of the data sets. The second is the capability to use a variety of visualization strategies, within the chosen coordinate system, for examining a single set of parameters from one source or many parameters from multiple sources. Specific representation techniques illustrate different aspects of data. Hence, no single method is always suitable, and not all techniques are useful for all data sets.

Therefore, methods that support the registration of multiple data sets in geographic coordinates, using similar cartographic warping of the respective data locations for the data sets in question, show promise as an alternative to meshing, interpolating, or resampling the original grids of each data set to a common rectilinear grid in projected space for correlative visualization in the Earth and space sciences. The former is used via the IBM Visualization Data Explorer developed by Visualization Systems, IBM Thomas J. Watson Research Center [Lucas et al., 1992]. This is in contrast to the latter approach used by the author via the NSSDC Graphics System, developed by NASA Goddard Space Flight Center [Treinish, 1989; Treinish and Goettsche, 1991].

2. STRATOSPHERIC OZONE DEPLETION AND THE POLAR VORTEX

The aforementioned approach for the analysis of multiple data sets via correlative visualization can be illustrated by an example scientific problem. There is a phenomenon known as the polar

vortex that occurs in the Earth's upper atmosphere (primarily the stratosphere) above Antarctica during the winter and early spring of every year [Schoeberl and Hartmann, 1991]. This effect is characterized by a cyclonic circulation pattern around the South Pole. Many researchers believe that ozone-destroying chemicals are trapped in this vortex during the cold and dark of Antarctic winter. Once spring begins and the polar region emerges from the long night, it is theorized that these substances react photochemically with ozone to break the molecule apart and thus aid in the creation of the so-called Antarctic ozone hole. Hence, in late winter, regions of ozone depletion around the pole begin to form. Within a few weeks the ozone hole is completely established. By late spring the vortex weakens, causing the ozone depletion region to fragment and eventually dissipate. The questions of interest are: what are the characteristics of the south polar vortex that can be derived from diurnal observations of atmospheric dynamics, and how do they relate to independent measurements of ozone? The study of the appropriate data sets for the Southern Hemisphere winter and spring (June through December) are relevant. The examination of a single year, 1987, is made because that year showed the greatest amount of ozone depletion until recent years [Krueger et al., 1992].

2.1 Total Column Ozone Data

Perhaps the most critical effort to study stratospheric ozone has been via observations made by the Total Ozone Mapping Spectrometer (TOMS) aboard NASA's Nimbus-7 spacecraft. Nimbus-7 is in a (polar) sun-synchronous orbit, which means that it can roughly provide global coverage of the Earth for its suite of instruments once per day. Each portion of the Earth was observed nominally under the same illumination conditions from day to day. Measurements made by Nimbus-7 TOMS show the daily global distribution of stratospheric ozone from late 1978 until early May 1993. It measured the total column density of stratospheric ozone by observing back-scattered solar ultraviolet radiation in seven spectral bands. Approximately 200,000 such measurements were made each day, which covered the entire globe [Fleig et al., 1986].

TOMS required sunlight to operate. Hence,

there are periods of missing data because of local polar winters (darkness) in addition to the usual data dropout problems associated with spacecraft observations. These regions are visible as gaps in various realizations of the data. They are *not* the ozone hole. The data have been gridded in a regular lattice of 180 (1.0° in latitude) \times 288 (1.25° in longitude) from the raw observations for daily global coverage with cells, without data being flagged. The locations of missing cells are considered in all realizations. The total ozone measurements are in Dobson units (DU's), corresponding to a column density of 2.69×10^{16} molecules of ozone cm^{-2} .

Figures 3a and b show traditional two-dimensional visualizations of the ozone data on October 1, 1987, which is during the ozone depletion season. The rectangular presentation of the data is consistent with the provided mesh in that it is torn at the poles and at a nominal International Date Line. This cartographic representation of the Earth is known as a cylindrical equidistant or plate carré projection. The ozone data are overlaid with a map of world coastlines and national boundaries. In Figure 3a, the data are realized as iso-contour lines at 50 DU intervals, which indicate the spatial distribution of discrete thresholds within continuous data. In Figure 3b, a pseudo-color spectrum is used, which is linearly mapped over a range (110 to 650 DU) that is valid for the year of study (to provide consistent comparisons between single days or for animation), and it should provide a continuous representation of continuous data. Figure 3b also has a fiducial overlay (lines of latitude or parallels and longitude or meridians at 30° spacing) in white, which have been registered in this same rectilinear coordinate system. The grid cells where there are no data are visible as gaps in the pseudo-color realization. The area of low ozone is visible as a bluish band stretched across the bottom of the pseudo-colored rectangle.

Non-rectilinear cartographic projections are used relatively often by Earth scientists as a way of preserving area in a display of the entire globe compared to the cylindrical equidistant projection. Other projections may preserve shape or linear distance, for example, on selected portions of the globe [Pearson, 1990].

To examine a continuous phenomenon with a

Total Column Ozone (Dobson Units) -- October 1, 1987




Figure 3a. Contour-mapped global column ozone on October 1, 1987.




Figure 3b. Pseudo-color-mapped global column ozone on October 1, 1987.

central focus far from the Equator and the Prime Meridian, such as the ozone hole, a non-cylindrical projection is required. Figure 4 illustrates the same data as in Figure 3, except polar orthographic projections for both the Southern and Northern Hemispheres are employed, which are shown in the left and right side, respectively, of the figure. For the orthographic projection, all meridians are straight lines radiating from the central pole. The

parallels are concentric circles, which become compressed toward the Equator. The orthographic projection can be characterized by

$$x = R \cos(\varphi) \cos(\theta) \quad (1)$$

$$y = R \cos(\varphi) \sin(\theta)$$

where $[\varphi, \theta]$ represents the location of each node




Figure 4. Polar orthographic-warped pseudo-colored deformed surfaces of total ozone.

on the Earth's surface in the original mesh as $[\text{latitude}, \text{longitude}]$, R is a scaling radius (e.g., 90°) and $[x, y]$ represents the location of each node in the deformed, curvilinear (orthographic) coordinate system with an assumed pole point of 90° north latitude and 0° east longitude and an assumed pole point of 90° south latitude and 0° east longitude for the Northern and Southern Hemispheres, respectively.

In addition to the pseudo-color spectrum, this two-dimensional cartographic projection of the data is extended by redundant realization as a deformed surface (i.e., both height and color correspond to ozone density). The bluish contiguous area over Antarctica clearly depicts the depletion region, illustrating the advantage of choosing an appropriate cartographic coordinate system. The height mapping clearly dramatizes the concept of a depression in the ozone layer, while the color enhances this perception as color enhances a topographic map. It provides a continuous representation consistent with the spatial structure of the data and, in animation, presents the dynamics in an easily

discernible fashion.

The use of hue-based pseudo-color mapping for realization and rendering of data as images can create problems in interpretation because of how the human visual system responds to color [Rogowitz et al., 1992]. For example, discontinuities appear in the pseudo-color representation, such as in Figure 3b, that are not in the data but are caused by how humans perceive hue (i.e., discretely). However, such pseudo-color maps are virtually standard in many disciplines. The acceptance of alternate, perceptually correct pseudo-color maps (e.g., luminance-based [Lefkowitz and Herman, 1992]) would be limited in these disciplines because of their unfamiliarity. Therefore, the introduction of redundant realization techniques, such as surface deformation, retains the familiar pseudo-color scale but helps lessen its negative perceptual impact.

Below each translucent surface is a hemispherical map that has been registered in the same orthographic coordinate system as the ozone data. This map consists of a monochromatic topographic

surface, which is deformed based upon height above or below sea level. The monochromatic scale is chosen to impart the appearance of a topographic map (e.g., the oceans are dark gray) for each hemisphere. This surface is created from a topographic data base on a rectilinear grid at 0.5° resolution. The grid is warped by the same orthographic projection that was applied to the ozone data, although the data were originally in a different coordinate system at different resolution. Both the ozone and topographic surfaces are Gouraud shaded. In addition, the topographic map is overlaid with the same coastline map in magenta, with political boundaries corresponding to each hemisphere, that was used in Figures 3a and b. The map geometry was transformed in a manner similar to that of the ozone and topographic data.

There is a seam in each of the orthographic surfaces corresponding to where east longitude is either -180° or $+180^\circ$, nominally the International Date Line, which is an artifact of the warping of the original rectilinear data onto a continuous surface, welding the discontinuity in the provided form of the data. The use of coordinate warping does preserve this inherent discontinuity in the data, which would not be the case if traditional interpolation techniques were chosen. In general, this seam will not smoothly connect the surface because of how TOMS gathers data. This scanning instrument examined each portion of the Earth at a different time of day, still covering the entire globe once per day. Hence, observations on each side of that line were taken approximately 24 hours apart and usually are not the same.

2.2 Dynamics Data: Atmospheric Temperature and Winds

Global atmospheric dynamics data (such as temperature and wind velocity) are often derived from spacecraft, balloon, and aircraft observations, which have been modeled and gridded on a 2.5° grid, 144×73 cells (longitude \times latitude) at different levels in the atmosphere, based on their pressure. Hence, a two-dimensional slice of these data at a specific pressure level is organized in a torn mesh similar to that of the total column ozone, but at lower resolution and in a different geographic coordinate system. These data may also have gaps

in coverage, including only a partial value for some wind cells (i.e., one or two of the three vector components are missing).

If this cartographic theme is extended to three dimensions by performing a Cartesian to spherical coordinates transformation, then for a "spherical" projection, all meridians are great circles converging at the poles. The parallels are also great circles, which become compressed toward the poles. The spherical projection can be characterized by

$$\begin{aligned} x &= (h + r) \cos(\varphi) \sin(\theta) \\ y &= -(h + r) \cos(\varphi) \cos(\theta) \\ z &= (h + r) \sin(\varphi) \end{aligned} \quad (2)$$

where $[\varphi, \theta, r]$ represents the location of each node on the Earth's surface as [latitude, longitude] at its radial distance $[r]$ from the Earth's center in the original mesh, h is the height above the surface of the Earth, and $[x, y, z]$ represents the location of each node in the deformed, curvilinear (spherical) coordinate system.

For the data being examined, there are seven pressure levels (1,000, 850, 700, 500, 300, 200 and 100 millibar [mb]). The data are treated as a true volume by warping the parallelepiped mesh representation of the atmosphere into a collection of concentric spherical shells that compose a volume of 73,584 nodes. The maps used in Figures 3 and 4 are replaced by a topographic globe such that all data are registered in a common, spherical, Earth-centered coordinate system. The radius of the globe is chosen to be the same as the inner radius of the spherical shell corresponding to the 1,000-mb level.

Figure 5 illustrates these ideas for both the temperature and wind data on October 1, 1987. For the temperature data, surface extraction techniques are employed because direct-volume rendering shows little quantitative information. The temperature data are realized as pseudo-color and opacity-mapped isothermal surfaces. The isosurfaces are at 194 and 294 K with the higher value (closer to the Earth's surface) being more opaque. The pseudo-color spectrum on the left corresponds to that of the temperature isosurfaces. The wind data are realized via streamlines generated by numerically integrating the paths taken by injecting massless particles in the wind field at 150 points uniformly distributed within the volume. The lines are




Figure 5. Isosurfaces of atmospheric temperatures and streamlines of atmospheric wind velocity.




Figure 6. Temperature surface (194 K) pseudo-colored by pressure height.

pseudo-color mapped by horizontal speeds ranging from 0 to 80 m/sec, which correspond to the pseudo-color spectrum on the right. The vertical component of the wind ranges from about -33 to +45 mb/sec.

Further study of the data using these techniques shows that there is a region of cold air (around 200 K) over the Antarctic during this period that is concentrated near 100 mb in pressure height, and that is surrounded by high-speed winds. The shape of this cold air mass and its diurnal variation at first glance seem to be similar to the depletion region in the total column ozone data. The shape of this air mass can be seen in Figure 6. Two spherical, wire-frame meshes surround a monochromatic globe. The outer green mesh of quads is at the 100-mb pressure level. The inner blue mesh of triangles is at the 200-mb pressure level. An isosurface at 194 K is shown for October 1, 1987, which has two components, south polar and equatorial.

The isosurface is pseudo-colored on a linear scale according to the pressure height at which the 194 K value is determined. The pseudo-color scale

ranges from blue (100 mb) to green (200 mb). Regions below 200 mb are colored pink. The south polar component of the isosurface is mostly at the 100-mb level. Hence, further study focuses on 100-mb data.

2.3 Correlating Ozone With Temperature and Winds

Two different approaches to visually correlating the ozone and dynamics data for the 100-mb level are taken applying the concept of coordinate warping to achieve geographic registration. The 100-mb data are the same as those used in Figures 5 and 6; hence, they are on a different grid than that of the ozone data. Because this study examines a phenomenon that is focused on a polar region and is nearly hemispheric in geographic extent, the orthographic cartographic projections discussed in equation (1) are used. They are illustrated using data from October 1, 1987, in an attempt to show the formation of the polar vortex and the ozone hole itself in Figures 7 and 8.

Figure 7 shows four separate and different data-driven representations of the atmosphere over the




Figure 7. Southern Hemisphere Orthographic Warped Atmospheric Data on October 1, 1987.

Southern Hemisphere in the same geographic coordinate system using the same three redundant realization techniques: pseudo-color mapping, surface deformation, and iso-contouring. This is the same approach as used for the ozone data in Figure 4, except for the addition of contour lines. In the upper left is the ozone column density with contours every 50 DU, from 100 to 650 DU. The upper right shows the 100-mb temperature with contours every 5 K, from 180 to 245 K. The lower left shows 100-mb horizontal wind speed with contours every 10 m/sec, from 0 to 80 m/sec. Cells where one or more components of the wind velocity are missing are shown as gaps in this surface.

In an attempt to show the correlation among these observable quantities in data space, a simple linear model is constructed such that

$$M = C_O Q + C_T T + C_V |\underline{v}| \quad (3)$$

where Q is the normalized total column ozone ranging from 0 to 1 (scaled for the dynamic range of 110 to 650 DU); T is the normalized 100-mb temperature ranging from 0 to 1 (scaled for the dynamic range of 180 to 245 K); $|\underline{v}|$ is the normalized 100-mb horizontal wind speed ranging from 0 to 1 (scaled for the dynamic range of 0 to 80 m/sec); M is a scalar field representing a unitless linear combination of the three parameters ranging from 0 to 3; and C_O , C_T , and C_V are normalized weighting factors, which are set to 1.

The ozone data are bilinearly interpolated to the grid on which the 100-mb data are available prior to the normalization. Independent gaps in both the ozone and wind measurements are properly maintained in the computation of M , which is shown in the lower right with contours every 0.5, from 0.0 to 3.0. Each of the surfaces shows a similar structure—a depression of comparable shape and areal extent over Antarctica for low ozone, temperature, and wind speed, respectively, each with a boundary corresponding to that of the polar vortex. The relative contribution of each of these parameters to the depression structure for M can be examined by interactively adjusting the weighting factors (C_O , C_T , and C_V) in the computation of M .

Figure 8 combines each of the parameters into one visual object. As with Figure 4, both hemi-

spheres are shown in the left and right sides, respectively, of the figure. The data are stacked vertically and shown with topographic, coastline, and national boundary maps. The difference between Figures 8 and 4 are representations for the 100-mb horizontal wind velocity and temperature stacked between that of the ozone and the maps. Below the ozone surfaces are plates of vector arrows representing horizontal winds, whose directions correspond to the direction of the wind and size and color correspond to wind speed, ranging from 0 to 80 m/sec.

Below the winds and above the maps are flat, translucent planes corresponding to the 100-mb temperature realized as pseudo-color-mapped, filled isothermal contours every 5 K over the range of 180 K to 245 K.

A daily animation sequence from June through September shows the availability of polar ozone data as well as the formation of the depletion region in both presentations. A similar signature is visible as a blue region in the temperature data, before the ozone depletion occurs, as a cold air mass forms over Antarctica in the winter and persists into spring. An analogous but less well-defined shape forms in the wind realization. With the onset of spring, ozone observations become available and form a similar depression pattern to that of the temperature and wind. During this period, daily animation shows a clear rotation of the depletion region surrounded by the boundary of the polar vortex.

This rotation, which has a period of several days, is synchronous between the 100-mb and ozone data. The arrangement of the wind velocity arrows in Figure 8 evoke a cyclonic pattern corresponding to the polar vortex, which appears almost steady-state in the winter and early spring. By early November, the warming of the upper atmosphere over Antarctica is obvious with direct correspondence to the dissipation of the polar vortex in the wind data and the breakup of the ozone depletion region.

Although the four time-varying surfaces do show the correlation among the ozone and dynamics data, they are difficult to observe, especially in animation. The eye tends to focus on one or two of the surfaces only. Therefore, at the cost of obscur-




Figure 8. Polar orthographic-warped atmospheric data on October 1, 1987.

ing some of the data, the stacked approach shows the synchronous circulation in each data set for the Southern Hemisphere during this period at a single glance.

3. IMPLEMENTATION

The techniques described herein have been developed via the IBM Visualization Data Explorer (DX), a general-purpose software package for scientific data visualization and analysis. It employs a data-flow-driven client-server execution model and is currently available on UNIX workstations (Sun, Silicon Graphics, Hewlett-Packard, IBM, and Data General), as well as a medium-grain, coherent shared memory parallel supercomputer, such as the IBM Power Visualization System [Lucas et al., 1992]. DX simplified the implementation of cartographic warping and its simultaneous application to disparate data sets for correlative visual display by providing an extensible tool kit of polymorphic operations that are interoperable and seem typeless

to the user. This polymorphism is a consequence of DX being built on a foundation of a unified data model, which describes and provides consistent access services for any data that are to be studied, independent of shape, rank, type, mesh structure, or dependency or aggregation. As a result, regular and irregular, structured and unstructured data are uniformly supported, as well as regions where data are missing.

The relevance of DX to correlative visualization problems is illustrated with a simple example. Although DX has several interfaces, the techniques discussed herein utilized visual programming, in which each computational task is assigned an icon and the flow of control and data are defined by connecting the icons as a direct acyclic graph. Figure 9 shows a visual program that generates a visualization of total column ozone as a radially deformed, pseudo-color- and opacity-mapped spherical surface surrounding a globe. Figure 9 also shows the pseudo-color and opacity map for the ozone surface via a color map editor and the

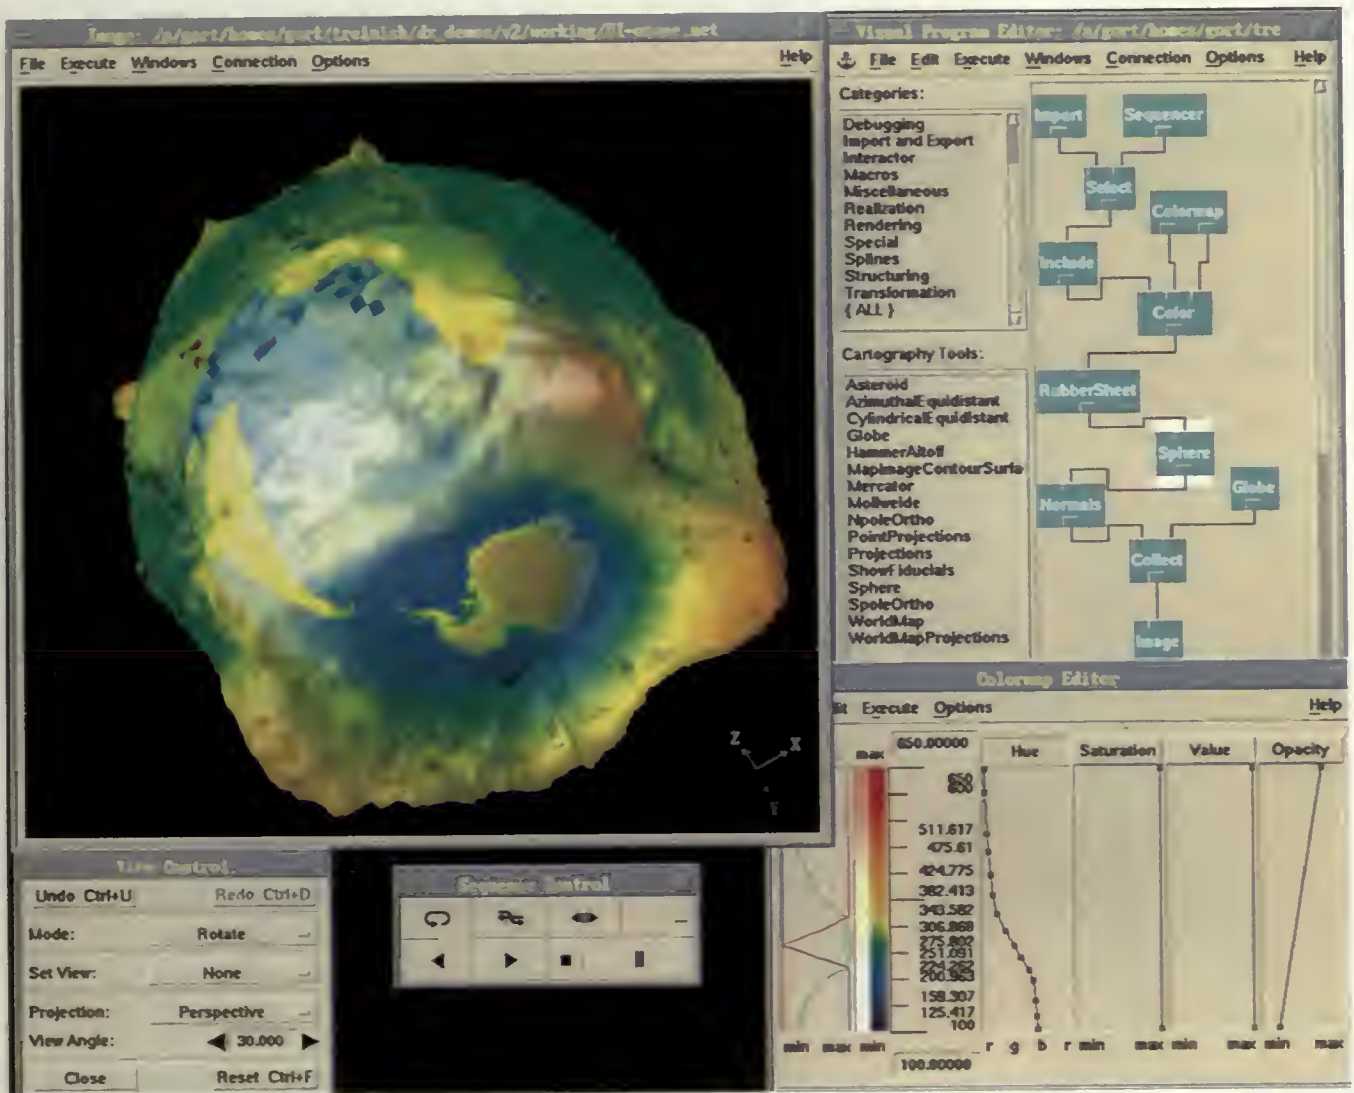


Figure 9. Example DX program to visualize global column ozone on October 1, 1987.

resultant image of the ozone data registered with a globe. Similar visual programs would show the atmospheric temperature around a globe, although the data are different in structure.

The program in Figure 9 has the following key operations:

1. *Import*: read data from the disk.
2. *Include*: subset data by value and indicate invalid data.
3. *Color*: assign pseudo-color and opacity maps.
4. *RubberSheet*: deform mesh by value.
5. *Sphere*: warp mesh onto a sphere.
6. *Normals*: compute normals for Gouraud shading.
7. *Globe*: generate a globe representation of the Earth.
8. *Collect*: aggregate inputs into a single entity.

9. *Image*: render an image and provide direct interaction.

For direct rendering of volumes, the *RubberSheet* operation would be eliminated. For extraction of surfaces from volumes, the *RubberSheet* module would be replaced with *Isosurface*. For a pseudo-color image (as in Figure 3b), steps 1, 2, and 3 would be used, optionally with a cartographic projection (e.g., *Mollweide*) and then *Image* (step 9). The *Sphere* and *Globe* operations are implemented as macros (visual programs without custom coding). The *Globe* operation uses steps 1, 3, 5, and 6 (with an option for 4) applied to topography data. The *Sequencer* operation is for the sequence control used as a virtual VCR for specifying animation. The *Select* operation

identifies a member of time series. The *Colormap* operation is for the color map editor used for the mapping of data values to hue, saturation, value, and opacity.

4. SUMMATION AND FUTURE WORK

The application of cartographic warping to the correlation of global atmospheric data sets yields visualizations that illustrate a simple notion about the possible relationship between temperature and winds, as well as their contribution below the tropopause to the formation of the polar vortex and ozone depletion. The use of interpolation prior to realization is unnecessary for the visualization of gridded data in a system with a sufficiently robust infrastructure of data structure and geometric support.

The ideas introduced with the analysis of observational data related to ozone and tropospheric dynamics can be extended by considering the correlation between these same ozone data and objective analyses that include the entire troposphere and stratosphere as a continuum. Visual correlation of these data is also useful for the examination of potential depletion regions in the Northern Hemisphere or the dynamics conditions for their possible formation. Data reflecting the typically disturbed conditions in northern polar regions in mid-winter through spring relative to the Southern Hemisphere yield more complex realizations than is the case for comparable southern polar regions. Although these data require more care in their presentation, an approach similar to that used for the Southern Hemisphere still applies. In addition, comparison of these data with spacecraft observations of clouds may yield additional insight into the dynamics of ozone depletion because polar stratospheric clouds are believed to provide sites for conversion of ozone-destroying chemicals from inactive forms to highly reactive ones during the darkness of polar winter [Hamill and Toon, 1991].

Acknowledgments. All of the data sets discussed herein were provided courtesy of the National Space Science Data Center at NASA Goddard Space Flight Center in Greenbelt, MD.

REFERENCES

- Fleig, A.J., P.K. Bhartia, C.G. Wellemeyer and D.S. Silberstein, Seven Years of Total Ozone From the TOMS Instrument—A Report on Data Quality. *Geophys. Res. Lett.*, 13, 12, November 1986.
- Hamill, P. and O.B. Toon, Polar Stratospheric Clouds and the Ozone Hole, *Physics Today*, 44, 12, 34-42, December 1991.
- Krueger, A., M. Schoeberl, P. Newman, and R. Stolarski, The 1991 Antarctic Ozone Hole: TOMS Observations, *Geophys. Res. Lett.*, 19, 12, 1215-1218, June 1992.
- Lefkowitz, H. and G.T. Herman, Color Scales for Image Data, *IEEE Computer Graphics and Applications*, 12, 1, 72-80, January 1992.
- Lucas, B., G.D. Abram, N.S. Collins, D.A. Epstein, D.L. Gresh, and K.P. McAuliffe, An Architecture for a Scientific Visualization System, *Proceedings of IEEE Visualization '92*, 107-113, October 1992.
- Pearson, Frederick, II, *Map Projections: Theory and Applications*, CRC Press, 372 pp., 1990.
- Rogowitz, B.E., D.T. Ling, and W.A. Kellogg, Task Dependence, Veridicality, and Pre-Attentive Vision: Taking Advantage of Perceptually-Rich Computer Environments. *Proceedings of SPIE Human Vision, Visual Processing, and Digital Display III*, 1666, 504-513, February 1992.
- Schoeberl, M.R. and D. Hartmann, The Dynamics of the Stratospheric Polar Vortex and its Relation to Springtime Ozone Depletions. *Science*, 251, 46-52, 1991.
- Treinisch, L.A., An Interactive, Discipline-Independent Data Visualization System. *Computers in Physics*, 3, n. 4, July 1989.
- Treinisch, L.A. and C. Goettsche, Correlative Visualization Techniques for Multidimensional Data, *IBM Journal of Research and Development*, 35, 1/2, January/March 1991.

LinkWinds: A Visual Data Analysis System and its Application to the Atmospheric Ozone Depletion Problem

ALLAN S. JACOBSON and ANDREW L. BERKIN
*Jet Propulsion Laboratory, California Institute of
 Technology, Pasadena, CA*

The Linked Windows Interactive Data System (LinkWinds) is a prototype visual data exploration system resulting from a NASA Jet Propulsion Laboratory (JPL) program of research into the application of graphical methods for rapidly accessing, displaying, and analyzing large multivariate multidisciplinary data sets. Running under UNIX, it is an integrated multi-application executing environment using a data-linking paradigm to dynamically interconnect and control multiple windows containing a variety of displays and manipulators. This paradigm, resulting in a system similar to a graphical spreadsheet, is not only a powerful method for organizing large amounts of data for analysis, but leads to a highly intuitive, easy-to-learn user interface. It provides great flexibility in rapidly interacting with large masses of complex data to detect trends, correlations, and anomalies. The system, containing an expanding suite of non-domain-specific applications, provides for the ingestion of a variety of data base formats and hard-copy output of all displays. Remote networked workstations running LinkWinds may be interconnected, providing a multi-user science environment (MUSE) for collaborative data exploration by a distributed science team. The system is being developed in close collaboration with investigators in a variety of science disciplines using both archived and real-time data. It is currently being used to support the Microwave Limb Sounder (MLS) in orbit aboard the Upper Atmosphere Research Satellite (UARS). This paper describes the application of LinkWinds to this data to rapidly detect features, such as the ozone hole configuration, and to analyze correlations between chemical constituents of the atmosphere.

1. INTRODUCTION

Recent advances in remote sensing capabilities and computational power are providing unprecedented ability to study our world. These improvements are also producing an ever-increasing flood of data that must be gathered, transported, stored, and analyzed to be fully utilized. This paper reports on a system capable of facilitating the rapid visual analysis of large masses of data, and discusses its application to upper atmospheric science. This system grew out of a NASA project at the Jet Propulsion Laboratory to study the application of computer graphics to the problems of quickly and interactively exploring and analyzing very large amounts of scientific data. The objectives of the program are to (1) develop a software environment that will support the rapid prototyping of visual data analysis applications, while at the same time maintaining the high level of performance necessary for interactively manipulating graphical displays; (2) develop a user interface that is truly intuitive, allowing quick access to the software for the novice as well as the advanced user; (3) provide a suite of sample applications that are useful across a variety of scientific disciplines; and (4) provide tools to support user development of applications for this environment.

2. LINKWINDS

The Linked Windows Interactive Data System, or LinkWinds, is a prototype product of this research effort. In compliance with the research objectives, LinkWinds, an integrated multi-application execution environment with a full graphical user interface (GUI), is a visual data analysis and exploration system designed to rapidly and interac-

tively investigate large multivariate and multidisciplinary data sets to detect trends, correlations, and anomalies. The system, operating under UNIX, is based on an object-oriented programming model and is implemented in the C language. It draws upon the Silicon Graphics, Inc. (SGI), GL library for its GUI and graphics support software and presently runs only on workstations supporting this library. This includes all SGI workstations and those of other manufacturers that have licensed and support the GL library. Third-party software and the advent of Open GL as a standard graphics library will greatly increase the portability of LinkWinds in the near future.

Data sets and individual tools for display or control of the data are coded as objects, each occupying a window on the LinkWinds screen and communicating with other objects through a message-passing protocol. The objects or windows can be linked or unlinked at the discretion of the user. Linking the windows sets up one-way message paths between objects. This data linking paradigm makes the system perform much like a graphics spreadsheet and, as in a spreadsheet, is a powerful way of organizing the data for analysis while providing a natural and intuitive interface. Data-linking, and its user interface implications, are discussed below.

Messages generated by LinkWinds objects are recorded as program statements in an underlying language called Lynx. The message-passing characteristics are the basis for two key LinkWinds functions. The first is the maintenance of an internal journal of all user-originated commands executed by the environment. This file can be saved at any time through a menu option. The record can then be replayed at the initiation of subsequent LinkWinds sessions, allowing the user to draw on a previous layout of LinkWinds applications and links or repeat a full analysis session.

The second function based on the Lynx message-passing protocol is the multi-user science environment (MUSE), which provides a method for multiple LinkWinds systems to communicate via networks. Using menu options, remotely separated users can connect to one another and, by establishing a telephone voice connection, can cooperatively view and manipulate their data. A successful connection requires that each user be executing

LinkWinds and that each has access to the data sets being analyzed. This is normally arranged by transporting the data sets prior to the collaborative session. Because only the messages are sent, and not the actual data, a very low bandwidth is required, making for quick and efficient communication. The MUSE capability is also used to provide tutorials over the network to new users and to allow users to demonstrate recommendations for application changes or to point out bugs.

Hard copy of the LinkWinds displays are provided by function keys on the keyboard. Placing the cursor in a window and pressing F1 will produce an image of a window's contents; pressing F2 will save the complete window and frame; and F3 will save the full screen. The figures shown were obtained in this manner.

3. DATA LINKING AND THE USER INTERFACE

In addition to the normal GUI functions provided by the windowing environment, dynamic manipulation of graphs and images is facilitated through the data-linking paradigm. Data linking can be understood in the context of a spreadsheet, where cells containing numbers are linked to other cells. Formulas are associated with each cell, so that when a number changes, all cells linked to the changed cell recalculate their values. LinkWinds does the same thing, but in a graphics environment where the rigid grid structure gives way to free form, and a cell can translate, for instance, into sliders or large-scale number arrays such as images.

This data-linking paradigm is one of the most distinguishing features of LinkWinds; it evolved from a desire to create a truly easy-to-learn and intuitive user interface. A guiding principle is that users are impatient and want to get started on productive work as quickly as possible. Large manuals only discourage them [M. Rettig, 1991]. Therefore, an interface was needed that can be learned by exploration and that conforms to expectations as the user works.

Data linking is affected through two icons. The link icon is a button displaying two interlocking rings, while the unlink icon displays two rings that are separated. Objects on the screen may have a single link button, the full set of link and unlink buttons, or no buttons. The presence of a single link button indicates

a data object, while the presence of the pair indicates applications with control functions. A window with no buttons is an application with only display capabilities. To perform a link, the cursor is placed on the appropriate button, and a “rubber band” is dragged out and dropped into the application to be linked. To break the link, the same routine is followed using the unlink button. The rubber bands signify that the links may be either continuously displayed or hidden during the session, depending on the user’s preference. There are two simple rules to follow in applying the linking paradigm:

1. When, as a result of menu selections, an empty window appears on the screen, put data in it. This is done by linking a data object into the window.
2. When an object with the pair of link symbols appears, exercise its control function by linking it into any application object.

Our experiences with users has confirmed the intuitiveness of this paradigm. Scientists quickly become familiar with the manipulations that control LinkWinds, often grabbing the mouse out of our hands during demonstrations to try themselves. New users typically require about 30 minutes of demonstration to understand the system well enough to embark upon productive work. Recently, groups of incoming Caltech graduate students were sent for a LinkWinds tutorial, and they quickly were performing competently.

While the LinkWinds scheme of linking together objects on the screen is reminiscent of dataflow systems, the similarity is only superficial. LinkWinds is a multi-application execution environment and therefore has significant architectural advantages. Any data transformations are done within high-level applications, which access the data through pointers to shared memory, negating the need to store intermediate copies of the data. This results in more efficient use of memory and better performance. LinkWinds is optimized for execution rather than programming and, with no need for linguistic completeness, is less complex and therefore more easily learned.

4. DATA BASE INTERFACE

The current version of LinkWinds interfaces with both archived and real-time data. In the

archived data mode, the primary data interface is with the widely accepted 8-bit raster and scientific Hierarchical Data Format (HDF), created and supported by the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign. The 8-bit raster format allows three-dimensional data files to be constructed as a sequence of images. LinkWinds also accepts data in a byte format, the Silicon Graphics, Inc. RGB image format, and the Common Data Format (CDF), originated and supported by the Goddard Space Flight Center (GSFC). Additional data formats are imported using DataHub [Handley and Rubin, 1993], a system being developed in close collaboration with the LinkWinds effort, and also supported by the NASA Applied Information Systems Research Program. DataHub provides LinkWinds direct access to a wide variety of data formats, including the Airborne Visible and Infrared Imaging Spectrometer (AVIRIS) format, Planetary Data System (PDS) format, and NetCDF, a derivative of CDF developed and supported by the National Center for Atmospheric Research (NCAR). The two programs are working together so that DataHub and LinkWinds may be called from inside the other, thus complementing their capabilities. LinkWinds will use DataHub for additional formats and to reduce large data sets, while DataHub will use LinkWinds for visual data selection, subsetting, and validation.

Data sets to be ingested by LinkWinds are listed in a text file and appear in the top-level “data bases” menu. Wild cards may be used in this text file, allowing whole classes of data files to be accessible. New data files created during a session by DataHub or LinkWinds’ subsetting tool may be automatically added to the menu. Metadata needed to translate data and axis values to meaningful numbers, such as the number of axes and their names, are obtained from data files in the standard formats. A special metadata file exists for the raw byte data or to provide additional information not contained in the standard formats.

The real-time mode of LinkWinds is a recent development [Jacobson and Berkin, 1993], and it must be exercised through the creation of a server tailored to the format of the input data stream. The data description file contains the name of the server, as well as the usual auxiliary metadata information.

The server is connected to LinkWinds through UNIX sockets. An interrupt system within LinkWinds senses the arrival of data and notifies the pertinent applications. In addition to the actual measured data, clock times and engineering information about the status of the sending devices may be read. Incoming data are saved in HDF 8-bit format and may then be replayed by requesting a replay server in the metadata file. Specific data files and time intervals may be chosen for replay. As a test demonstration, LinkWinds was used to monitor and analyze data from the University of Iowa's Plasma Wave Subsystem aboard the Galileo spacecraft during the Earth 2 encounter in December 1992.

5. APPLICATION TO ATMOSPHERIC DATA

A suite of applications useful across many disciplines has been developed for the LinkWinds environment (see Table 1). Figure 1 shows a typical session to explore a data set collected by the Microwave Limb Sounder (MLS) currently in orbit aboard the Upper Atmosphere Research Satellite (UARS) [Waters et al., 1993]. These data are three-dimensional, with the concentration of various chemical constituents varying with latitude, longitude, and pressure. The instrument has sensitivity over the range from 100 millibars (mb) to 0.1 mb. A new data set is collected daily; the data shown here are from day 56 of the mission, which is day 310 of 1991.

The LinkWinds top-level menu appears in the center left of Figure 1. From this menu, the data bases, tools, and system options are selected. Data objects, with their single link buttons, are just above the menu. In this case, the data displayed are ozone and water vapor. The window, entitled Image1, contains a slice of the ozone data at a pressure of 21.54 mb, as selected by Slider1, which is linked to it. Slider1 is also linked to Image2, which shows the water vapor at the same pressure. Slider1 permits the user to scan the full data set from the maximum to minimum altitudes. The user can also switch to any of the three orthogonal axes and similarly scan them. The amount of each constituent is given by color, as indicated on the color bars at the bottom of the images, with red denoting high concentration and

blue low. Gray indicates missing data, where MLS could not measure because of the orbital position, orientation, and sensing range of UARS.

The Southern Hemisphere ozone hole is clearly seen at the lower left of Image1, and a high value of the water vapor is seen in the corresponding location of Image2. This anticorrelation is readily visible in Scatter1, where the points shown come from the region defined by a bounding box control embedded in Image and linked to Scatter1. This box, visible in the lower left of Image1, can be resized or moved as desired. For every point inside this box, Scatter1 plots the ozone vertically versus water vapor horizontally. In the lower text box of Scatter1, statistical quantities associated with the scattered data are given. Of particular note is the closeness of the correlation coefficient to -1, indicating high linear anticorrelation. Other information, such as third and fourth moments, linear best fit data of the scatter, the number of points scattered, and the chi-squared value, may be selected from the Scatter menu. The full range of statistical information or the points themselves may be saved to files via buttons.

Two other applications display data along one dimension. In the lower center of Figure 1, LinePlot1 displays ozone concentration as a function of pressure. The location in latitude and longitude of this plot is controlled by Image1, which has been linked into LinePlot1. The green plot corresponds to the frozen crosshair at the center of Image1, while the white plot corresponds to the red crosshair and is instantaneously updated as the crosshair is moved. The much lower white values reflect the paucity of ozone inside the hole. The Profile application to the left of LinePlot1 shows ozone concentration along the cyan line in the slice of Image1. This line may be interactively updated. The current Profile line starts in the ozone hole, goes through an ozone-rich region, and ends back in the hole. This is reflected in both the colors of Image1 as well as the height of the curve of Profile1, providing a different way of visualizing the data.

Because this data set is global, Figure 2 shows the ozone data displayed in Globe1 as both a color and height field rendered on a sphere. The ozone hole and adjacent regions of higher ozone are clearly seen. The water vapor data could have been




Figure 1. LinkWinds session to explore upper atmospheric ozone and water vapor measured by the Microwave Limb Sounder aboard UARS.




Figure 2. Examination of other features of upper atmospheric ozone measured by the Microwave Limb Sounder aboard UARS.

used for the height field, clearly demonstrating any correlations. Slider1 also sets the pressure slice of this display, and the height scale is determined by the rotary dial on the left side of the window. Pan-Zoom and 2-Axis Rotator controls linked to Globe1 allow it to be interactively positioned as desired.

The Combine tool allows mathematical manipulations to be performed on slices of data. A calculator called from Combine via a button contains the standard mathematical functions and allows the input of constants, as well as slices. The user can combine up to three slices of data, from the same set or mixed sets, in any mathematical expression desired. In Figure 2, these slices have been selected by linking in LinePlot, with the red, green, and blue sliders setting the levels as 21.54, 14.68, and 10 mb, respectively. The calculator has summed these three slices, and the resulting slice and histogram of the data distribution are displayed in Combine1.

Among the tools in Table 1, two deserving special comment facilitate the creation of animations for immediate display on the screen or subsequent recording on video or film. Frame Animator is frame based, with the user selecting starting and ending control values, and the number of frames desired in the animation. The other Animator is time based. The user sets any number of control positions, each with associated key times graphically selected by moving the hands on a clock. The Animator then interpolates between these set positions to easily make animations with great flexibility in the control. The desired frame rate is selected from the Animator menu. Rates include those for film, video, and a screen display mode, which makes the animation in real time.

6. FUTURE PLANS

Several developments are planned for the future to significantly improve the usefulness of LinkWinds. As in the past, we will continue to develop applications in collaboration with science users seeking to solve real problems. Where relevant, these applications will make use of modern rendering techniques that can be applied successfully in an interactive environment.

A major impediment to the use of any visualization tool is the difficulty users have in inputting

their data. Developments of LinkWinds and DataHub will continue to make this process as seamless and automatic as possible. A related issue is the types of data addressed by LinkWinds. Available tools for visual data analysis are generally confined to relatively well-behaved and rectangularly gridded data sets. LinkWinds' first venture from the common mold was into the realm of real-time data, creating a capability to ingest such data and build interactive applications for monitoring and analyzing it.

There are other major neglected categories of data that are quite common in scientific research and badly in need of tools to support their exploration and analysis. Several problem areas to be addressed in the future are (1) data sets in which there are significant sources of error, either statistical and/or systematic; (2) data sets that are ungridded samples, either sparse or numerous, from which the user desires to construct gridded data sets over extended regions; and (3) disparate-sized data sets from a variety of instruments that must be warped and/or co-registered for overlay or comparison.

In the future, we intend to pursue the development of a users' applications generator for LinkWinds. Currently, the layout of the objects, or widgets, in all of the windows is determined by a text file. The user can reconfigure these windows either by editing this file or interactively from a menu-selectable "redesign" mode. We intend to expand this tool kit approach to further allow users to throw widgets away, or add new widgets from a provided catalog. In conjunction, LinkWinds will generate a C code source module to make these widgets work. This code will be suitable for use as a template for the development of a full application. As experience is gained with this approach, and a planned conversion of the code to C++ is accomplished, we anticipate that the rendering and display processes will also lend themselves to a limited catalog of processes selectable by the user.

Aware that the only way to develop useful tools is in conjunction with research on meaningful problems, our policy has been to encourage users and potential users to contact us concerning LinkWinds' changes and needs. We have responded and will continue to do so to the limit of our resources. LinkWinds is currently in use at

Table 1. Current Suite of LinkWinds Applications.

Control	
Slider	Controls which slice of data is displayed along any of three orthogonal axes
3-Axis Slider	Controls the slice of data displayed along the three orthogonal axes (three sliders)
3-Axis Rotator	Rotates 3-D applications along three orthogonal axes (three sliders)
2-Axis Rotator	Moving in a 2-D area, rotates 3-D applications (one slider)
Pan-Zoom Slider	Controls 3-D applications by changing the viewpoint
Combine Slider	Determines three slices of data and their three constant offsets (six sliders)
Animator	Provides time-based animations of any application with an arbitrary number of key frames, with a variety of frame rates
Frame Animator	Provides frame-based animations of any application, with only the start and stop control values being set
Color Tool	Involves interactive data set palette manipulation allowing color editing, data ranges being ramped in color, or substitution of pre-defined palettes
Display/Control	
Combine	Combines up to three slices of data using standard mathematical functions entered in an embedded calculator
Compare	Compares the functional behavior of each point in a data set with a reference point using a variety of mathematical functions
Data Subset	Allows the user to interactively save portions of the displayed data into HDF
LinePlot	Plots the values along a straight line going completely through a data set parallel to any axis, and also functions as a slider to select three slices
Histogram	Displays the distribution of values in the 256 data channels for up to three slices, and provides filtering and color stretching
Image	Displays a single slice of data or a composite RGB image of three slices with embedded crosshair, bounding box, and line controls
Display	
Plane	Polygonally renders an image in perspective relief, with an optional accompanying height field, of either a single slice or RGB three-slice composite
Globe	Polygonally renders an image on a globe, with an optional accompanying height field, of either a single slice or RGB three-slice composite
Orthoview	Displays in a 3-D point-by-point rendering all the values in a data set between two limits
Profile	Displays the data values along a line drawn on the Image, Combine, or Compare tools
2-D Scatter Plot	At every location in a slice, plots the values of one data set against the other to show the correlation
3-D Scatter Plot	At every location in a slice, plots the values of three data sets in 3-D to show their correlation
TrackPixel	Gives numerical information about the data in Image, Combine, or Compare, both at a point and averaged over a bounding box
Real Time	
StreamPlot	Display data (through a strip-chart recorder) as either color or line plots, as a function of channel number vertically and time horizontally
StreamLine	Plots data value versus channel number, updating in real time with options of saving a spectrum and averaging in time
StreamPlane	Functions much like StreamPlot, but with the data polygonally rendered in relief using the data values for height and color
StreamClock	Provides timing information for the current data, giving date, day of year, time of day, and internal spacecraft time
ChannelSlider	Controls the range of channels to be viewed, allowing concentration on features of interest

more than ten institutions, being applied to problems in remote-sensed and field geology, atmospheric physics and chemistry, meteorology, oceanography, chemical spectroscopy, space plasmas, genetics, and cellular biology. As it evolves, we expect to interact with a wider distribution of scientists engaged in research spanning additional scientific disciplines.

Acknowledgments. In addition to the authors, the current LinkWinds development team members include Martin Orton and Mark Rubin at JPL and Philip Mercurio and James McLeod at the San Diego Supercomputer Center. Alumni who made important contributions are Brian Beckman, Leo Blume, Bonnie Boyd, and Mitch Wade. Lee Elson provided useful discussions on upper atmospheric science. Andrew Berkin acknowledges support from the National Research Council. We extend our thanks to Joe Bredekamp and his staff at NASA, whose sponsorship makes this work possible. The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA.

REFERENCES

- Handley, T.H. Jr. and M.R. Rubin, DataHub: Intelligent Distributed Science Data Management to Support Interactive Exploratory Analysis, *Computing in Aerospace*, 9, 91-101, AIAA, Washington, DC, 1993.
- Jacobson, A.S. and A.L. Berkin, LinkWinds, A System for Interactive Scientific Data Analysis and Visualization, *High Performance Computing 1993*, edited by A. Tentner, 214-219, The Society for Computer Simulation, San Diego, CA, 1993.
- Rettig, M., Nobody Reads Documentation, *CACM*, 34, 19-24, 1991.
- Waters, J.W., L. Froidevaux, W.G. Read, G.L. Manney, L.S. Elson, D.A. Flower, R.F. Jarnot, and R.S. Harwood, Stratospheric ClO and Ozone From the Microwave Limb Sounder on the Upper Atmosphere Research Satellite, *Nature*, 362, 597-602, 1993.

The Use of LinkWinds for the Validation and Analysis of 14 Years of Microwave Sounder Unit Daily Global Temperature Anomaly Data

MICHAEL E. BOTTS

Earth System Science Laboratory, University of Alabama, Huntsville, AL

ROY W. SPENCER

Earth Science and Applications Division, NASA Marshall Space Flight Center, Huntsville, AL

Temperature data derived from the Microwave Sounder Unit (MSU) provides an opportunity for investigating atmospheric temperatures on a global scale since 1979. Fourteen years of global data sets of daily temperature anomalies within the lower stratosphere and lower troposphere are being generated at NASA Marshall Space Flight Center.

LinkWinds, a visualization/analysis package under development at NASA Jet Propulsion Laboratory, has been extremely useful for validating and analyzing these data sets. LinkWinds provides the ability to interactively scroll and animate through the 10,220 images of temporal data, to selectively slice and view the data along latitude, longitude, or temporal axes, to interactively analyze spatial and temporal variability within the data, and to perform correlative analysis between various elements of the data. These capabilities have been invaluable in allowing the recognition of processing artifacts, as well as the effects that physical phenomena, such as the El Niños effects and the Mt. Pinatubo eruption, have had on atmospheric temperatures.

1. INTRODUCTION

The availability of global data sets extending over relatively long periods is vital to global change research. With regard to atmospheric temperature measurements and the unresolved issue of global warming, the most temporally extensive data sets are, unfortunately, derived from sparse ground stations that are often located near rapidly growing urban areas or other areas greatly affected by localized conditions. Furthermore, surface-based temperature measurements in oceanic or other remote regions are sparse in both temporal and spatial domains. Previous studies have shown that the Microwave Sounder Unit (MSU) sensor, located on several National Oceanic and Atmospheric Administration (NOAA) satellites since 1979, provides an excellent measure of atmospheric temperature variability for the entire Earth [Spencer et al., 1990].

Methods have been established for consolidating 14 years of data from several MSU sensors and for gridding and calculating global atmospheric temperature anomalies for the lower troposphere and lower stratosphere [Spencer et al., 1991; Spencer and Christy, 1992a, 1992b]. Data sets of limited temporal resolution and range were generated during the previous studies to validate the techniques against ground-based and radiosonde data.

At Marshall Space Flight Center (MSFC), daily atmospheric temperature anomaly data sets have been generated from the MSU sensor data, for the 14-year period extending from January 1979 to the present, for both the lower troposphere and lower stratosphere. These data sets are undergoing validation and quality control at MSFC, before being made available for distribution to global

change researchers. Although some validation and quality control have been accomplished using established data-processing procedures, The Linked Windows Interactive Data System (LinkWinds) is being used extensively to provide additional quality control through interactive visualization and analysis of the entire data set.

2. THE DATA

The present MSU temperature anomaly data sets consist of daily global images for 14 years from 1979 to 1992, for both the lower troposphere and lower stratosphere. This calculates to 10,220 images. The temperature anomalies are mapped to a 2.5° grid, resulting in an image resolution of 144 x 72. At present, the data sets are stored in the 8-bit image Hierarchical Data Format (HDF), with 1 year (365 to 366 images) of data per HDF file. The stratospheric and tropospheric anomalies are divided into separate files. Each HDF file is about 3.8 MB, or more than 100 MB of total data for the entire 14-year period.

The MSU daily temperature anomaly data sets are derived from a series of processing steps, each of which can be a potential source of error. These steps include:

- Calibration of raw MSU antenna brightness data
- Retrieval of atmospheric temperature from various MSU channels
- Navigation/coregistration
- Consolidation of data from multiple sensors (i.e., adjusting intersatellite calibration differences)
- Limb correction (i.e., off-nadir corrections)
- Calculation of an average daily temperature for a given location, based on the entire data set
- Daily averaging and spatial gridding

It is also known that several natural factors can cause artifacts within the data, including the presence of short-lived isolated thunderstorm cells, which can cause a bias in the daily average temperature measurement, and high surface elevations, which can result in erroneous tropospheric temperature measurements caused by interference from ground temperatures.

3. THE TOOL

LinkWinds is a visualization/analysis tool under development at NASA Jet Propulsion Laboratory (JPL) [Jacobson, 1992; Jacobson and Berkin, 1993a, 1993b]. LinkWinds provides an intuitive and highly interactive interface for data exploration and analysis, as well as for intercomparison and correlation of multiple data sets. LinkWinds is best suited for exploring or comparing “datacubes” or “stacked” data, which might consist of two spatial dimensions (such as latitude and longitude) and a third dimension of either spectral bands, time, a collection of variables, or the third spatial dimension (such as altitude or depth). However, LinkWinds does not restrict the user’s choice of axes, so that any two- or three-dimensional data can be analyzed. Furthermore, LinkWinds allows correlative analysis between two or more similarly configured datacubes, in addition to extensive exploration and analysis within a single datacube.

In the case of the MSU data sets, LinkWinds provides the ability to interactively scroll through the 10,220 images of temporal data, to selectively slice and view the data along latitude, longitude, or temporal axes, to interactively view and analyze spatial and temporal variability for a given position or region of interest, and to combine and correlate various elements of the data. Rudimentary data base management is accomplished in LinkWinds through the use of a user-defined hierarchical menu structure. The user is provided a selection of tools into which these data can be loaded, including:

- **Image display tools** for two-dimensional and three-dimensional data display
- **Data feedback tools** for line plots, histogram display, and pixel and bounding box tracking
- **Tools for intercomparing data sets**, such as two-dimensional and three-dimensional scatter plots, and compare and combine tools
- **Frame-based and time-based animation controllers**
- **Control widgets**, such as sliders, rotators, pan-zoom controller, color table editor, and data subsetting tool

Many of the display tools also have associated internal controllers that can be linked to other tools.

For instance, the *histogram* tool allows the user to perform rudimentary image enhancement within *image*, *globe*, and *plane* tools, while *line plot* allows the selection of components for a three-band (RGB) composite image.

Controlling functions can be linked to more than one tool, allowing simultaneous control over the analysis and display of two or more data sets. Figure 1 shows an example of the MSU data set and some of the tools available within LinkWinds, including the *image*, *globe*, *histogram*, *line plot*, and *combine* with calculator tools.

4. DATA SET VALIDATION AND ANALYSIS

Although some validation and quality control of the data set can be achieved by analytical methods, the interactive visual exploration of the data provided by LinkWinds has proved invaluable for

rapidly spotting erroneous data or artifacts beyond those found by analytical means alone. Having the ability to very rapidly scroll through a year or more of daily images has been useful for locating gross errors, as well as errors that may not be apparent on single images, but that become noticeable during image animation. The investigation of discovered artifacts was able to be refined and new sources of error discovered using several additional capabilities of LinkWinds, including in particular the ability to slice and scroll the data sets along any orthogonal axes, to generate line plots along these axes, and to retrieve additional feedback through the *histogram* and *pixel tracking* tools. In addition, an initial investigation of the relationship among various parameters was possible using the *scatterplot*, *combine*, and *compare* tools.

These capabilities have been invaluable in allowing the recognition of processing artifacts, as



Figure 1. An example of a few of the tools available within LinkWinds, including *image*, *globe*, *histogram*, *line plot*, and *combine* (with associated image calculator).

well as the effects that physical phenomena have on atmospheric temperatures. The features and artifacts observed using LinkWinds included:

- Validation/quality control
 - Valid data ranges/bad pixels
 - Surface elevation effects (mountains)
 - Misplaced orbits
 - Striping caused by improper limb correction
 - Bug in software where arrays not cleared
 - Thunderstorm signatures
- Analysis
 - El Niños effects
 - Mount Pinatubo eruption effects
 - Quasi-biennial oscillations
 - Strong inverse correlation between tropospheric and stratospheric anomalies in summer hemispheres

The initial stage of validation of the data set using LinkWinds was to determine the range of

valid temperature anomaly values. The initial range of the data extended between about -40 to $+40^{\circ}\text{C}$. LinkWinds provided the ability to interactively examine the validity of pixels in the lower and upper ranges of the data. As illustrated in Figure 2, the color table editor was used to set the color of the “bad pixel” value to a very conspicuous color—in this case, bright green. Setting the sliders on the *histogram* tool to various upper and lower limits, pixels that exceed these values took on the “bad pixel” color and could easily be distinguished from other pixels. It was then possible to quickly slide through the sequences of images and determine at what upper and lower values the pixels seemed to be valid, rather than spurious, as was the case in Figure 2. Using this procedure, it was determined that the valid range for these temperature anomaly data was between -20 and $+20^{\circ}\text{C}$ for the lower troposphere and between -25 and $+25^{\circ}\text{C}$ for the lower stratosphere.




Figure 2. Example of “bad pixel” values determined using LinkWinds image, histogram, color, and pixel tracking tools.

Ground temperature contamination of tropospheric measurements in areas of high surface elevation is difficult to recognize within a single image, but is easily recognized by two methods within LinkWinds. First, regions of consistent spurious values often become recognizable while scrolling or animating through the temporal sequences of images. This method is difficult to illustrate within a publication. The second method takes advantage of the ability to view and scroll data along various axes within LinkWinds. As illustrated in Figure 3, the *image* tool can be set to view a slice of the data along a latitude-time plane by means of a pop-up menu. A slider then can be set to scroll through longitude space until spurious linear discontinuities are spotted parallel to the time axis. In Figure 3, the days run from 1 to 365 along

the horizontal axes, and latitude along the vertical axis. A discontinuity is shown, which persists throughout the year at 20.9S, 66.7W, and it results from the presence of the Andes Mountains. The 5-day periodicity of the microwave signal, as determined using the *line plot* tool, is apparently related to changes in the satellite viewing angle from nadir to oblique.

Other artifacts that were discovered in the initial MSU data sets included misplaced orbits and striping caused by improper limb corrections, shown in Figure 4a and Figure 4b, respectively, as well as instances where the data array had not been set to zero for areas of missing data. In addition, unreasonably high temperature gradients, indicative of the presence of local thunderstorms, were discovered and corrected. Many of the artifacts in the data sets

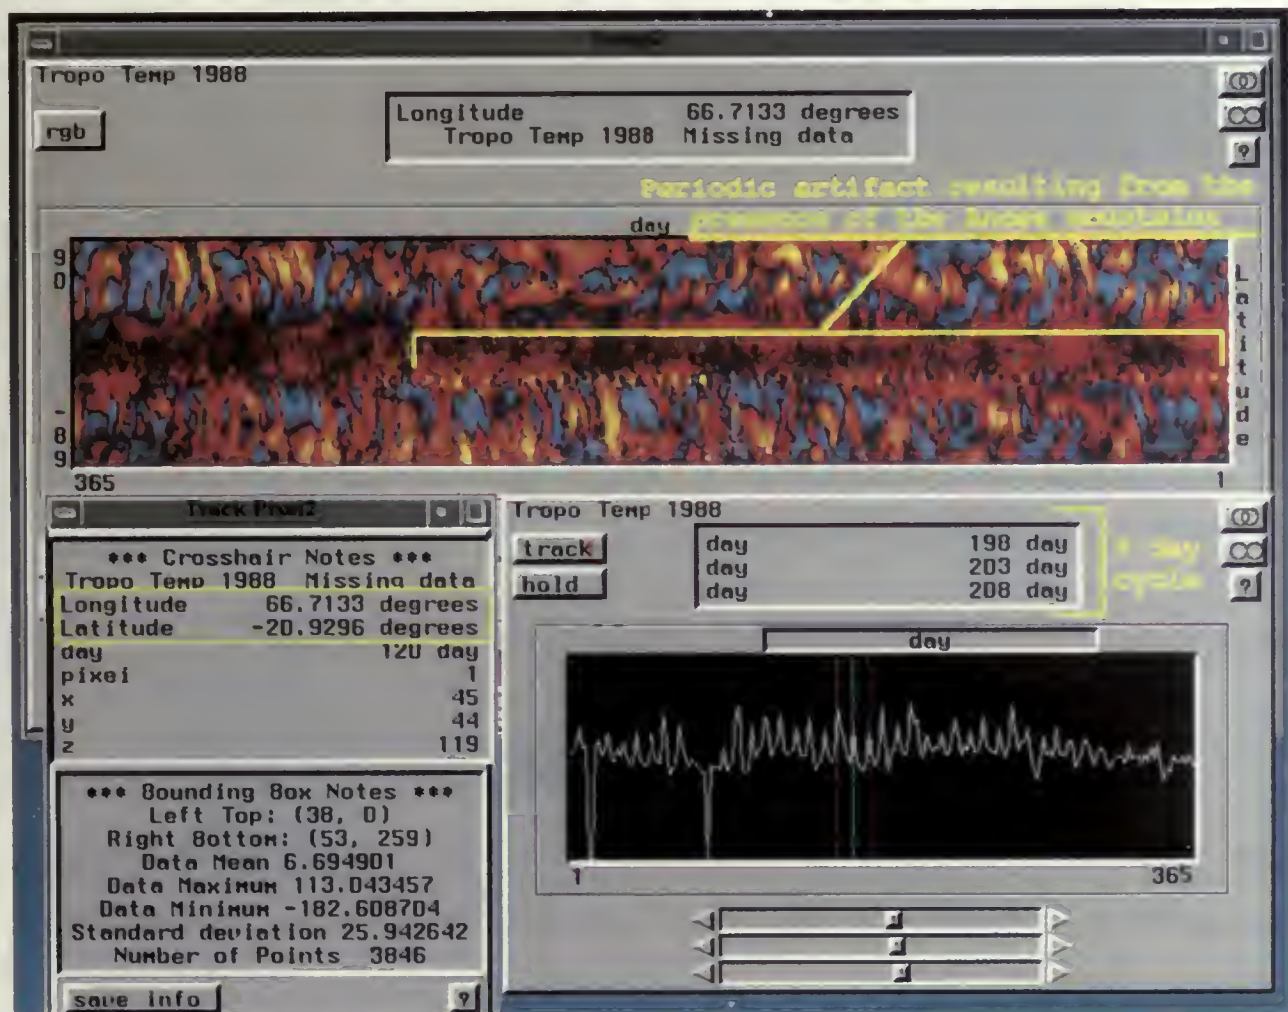


Figure 3. A latitude-day slice of the MSU tropospheric temperature anomaly data showing the presence of spurious data at 20.9S, 66.7W throughout the year, resulting from the presence of the Andes Mountains. The 5-day period of the artifact, as determined with the line plot tool, results from changes in the satellite viewing angle from near-nadir to oblique.




Figure 4. Additional artifacts discovered in the MSU temperature anomaly data using LinkWinds, including (a) misplaced orbits and (b) striping caused by improper limb correction.

would not have been recognized without the expertise of an atmospheric scientist, adding credence to the need for tools, such as LinkWinds, which are developed for and operated by the scientist.

In addition to recognizing undesirable artifacts, LinkWinds has also provided the opportunity to easily observe relationships between various natural phenomena and global atmospheric temperature anomalies. These phenomena include El Niño/southern oscillations (ENSO), quasi-biennial oscillations (QBO), volcanic eruptions, and a strong inverse correlation between stratospheric and tropospheric temperature anomalies within the summer hemisphere. Figure 5 shows the tropospheric warming in the tropics as a result of ENSO of 1983.

In Figure 6, longitude-day slices of the stratosphere, at 16.9N latitude, are shown for 1990, 1991, and 1992. The relative cooling of the lower stratosphere, possibly the indication of a QBO, can be seen in the second half of 1990 and the first half of 1991. As seen in Figure 6, this cooling cycle was interrupted by rapid stratospheric warming as a result of the Mt. Pinatubo eruption on June 27, 1991 (day 191).

An additional relationship that was observed with LinkWinds was a strong inverse correlation (correlation coefficient up to 0.93) between temperature anomalies in the lower troposphere and

those in the lower stratosphere within the summer hemisphere. In contrast, the stratospheric and tropospheric anomalies in the winter hemisphere consistently exhibit no correlation. Using the LinkWinds *scatterplot* tool, as shown in Figure 7, one can rapidly scroll through the year's worth of data and observe the daily scatterplot for the entire globe or for a user-selected subregion. Within each hemisphere, the stratospheric and tropospheric MSU temperature anomalies are observed to become highly inversely correlated as that hemisphere's summer months are approached, and then progressively degenerate to a condition of poor correlation as the winter months are approached. During the high correlation periods in the summer, the slope of the stratosphere/troposphere correlation curve is relatively consistent from year to year; it is around -1.45 ± 0.15 for the Northern Hemisphere and -1.35 ± 0.1 for the Southern Hemisphere. The intercept for the correlation curve varies from about 0 to 1.5°C for the Northern Hemisphere and from about -1.5 to 2.0°C for the Southern Hemisphere.

5. CONCLUSIONS

LinkWinds provides an intuitive and highly interactive interface and significant capabilities for validating and analyzing large data sets. Adequate quality control of 10,220 images within the



Figure 5. A latitude-day slice of the stratospheric and tropospheric MSU temperature anomalies showing the tropospheric warming in the tropics during the El Niño/southern oscillation (ENSO) in the first half of 1983.



Figure 6. Longitude-day slices at 16.9N latitude for stratospheric anomalies in the years 1990, 1991, and 1992. Note relative cooling from mid-1990 until mid-1991, possibly indicative of a quasi-biennial oscillation (QBO), which was suddenly interrupted by stratospheric warming resulting from the Mt. Pinatubo eruption of July 27, 1991.



Figure 7. Scatterplot of tropospheric versus stratospheric temperature anomalies during the summer in the Southern Hemisphere. Region of interest delineated by the bounding box in the troposphere image. Note correlation coefficient = -0.92 and other relevant information provided in the statistics box of the scatterplot tool.

MSU data set would not have been possible without LinkWinds, or a tool of its caliber. The ability to interactively compare and combine various parameters within a single tool makes LinkWinds an excellent tool for comparative analysis of multisensor and multiparameter data sets. LinkWinds also provides a promising development environment for future tools or for extension of the present LinkWinds functionality.

REFERENCES

- Jacobson, A.S., LinkWinds: An Approach to Visual Data Analysis, *Proceedings of Int'l Symposium on Spectral Science Research*, Maui, Hawaii, November 1992.
- Jacobson, A.S. and A.L. Berkin, LinkWinds, A System for Interactive Scientific Data Analysis and Visualization, *Proceedings of High Performance Computing Conference, Soc. for Computer Simulation*, 214, March 1993a.
- Jacobson, A.S. and A.L. Berkin, LinkWinds, A Visual Data Analysis System and its Application to the Atmospheric Ozone Depletion Problem, *Transactions of the AGU*, 1993b.
- Spencer, R.W. and J.R. Christy, Precision and Radiosonde Validation of Satellite Gridpoint Temperature Anomalies, Part II: A Tropospheric Retrieval and Trends During 1979-90, *J. Climate*, 5, 858-866, 1992a.
- Spencer, R.W. and J.R. Christy, Precision and Radiosonde Validation of Satellite Gridpoint Temperature Anomalies, Part I: MSU Channel 2, *J. Climate*, 5, 8, 847-857, 1992b.
- Spencer, R.W. and J.R. Christy, Precision Lower Stratospheric Temperature Monitoring With the MSU: Validation and Results 1979-91, *J. Climate* (in press), 1993.
- Spencer, R., J. Christy, and N. Grody, Global Atmospheric Temperature Monitoring with Satellite Microwave Measurements: Method and Results 1979-84, *J. Climate*, 3, 10, 1111-1128, 1990.
- Spencer, R., J. Christy, and N. Grody, Precision Tropospheric Temperature Monitoring 1979-1990, *Palaeogeography, Palaeoclimatology, Palaeoecology*, 90, 113-120, 1991.

Interactive Data Exploration and Particle Tracking for General Circulation Models

R.I. ROSENBAUM, R.L. PESKIN, S.S. WALTHER

Center for Computer Aids for Industrial Productivity,
Rutgers University, New Brunswick, NJ

H.P. ZINN

NASA Goddard Institute for Space Studies, New York, NY

The SCENE environment for interactive visualization of complex data sets is discussed. This environment is used to create tools for graphical exploration of atmospheric flow models. These tools may be extended by the user in a seamless manner, so that no programming is required. A module for accurately tracing field lines and particle trajectories in SCENE is presented. This is used to examine the flow field qualitatively with streamlines and pathlines and to identify critical points in the velocity field. The paper also describes a visualization tool for general circulation models on which the primary features of the environment are demonstrated.

1. INTRODUCTION

We have applied a distributed computing interface called SCENE (Scientific Computation Environment for Numerical Experimentation) to the analysis of atmospheric data sets. SCENE includes a "front-end" workstation, "back-end" computers, and software. The environment is designed to allow scientists and engineers to undertake numerical experimentation without programming, interact with their numerical experiments via interactive graphics, and have total access to the results through easy manipulation of a variety of graphical output structures. These structures are designed to enhance the user's ability to obtain quantitative information from graphics [Peskin et al., 1992].

The back-end systems support intensive numeric and symbolic computations, and this computational power is available to the user in a transparent manner. Thus, SCENE provides a way for the user to access high-performance computation without requiring direct involvement in complex aspects of computer programming systems. The environment is fully object oriented in design so that objects such as vectors and triangles in SCENE exist just as they do in physical models. This permits rapid and accurate prototyping of visualization tools for complex data sets. The environment employs a data base that recreates in memory the logical connectivity that exists in the grid of the data set. Maintaining a direct connection between the visual representation of data and the actual data set is the first priority of this data storage and retrieval model [Walther and Peskin, 1993].

1.1 Object Editor

To create a SCENE visualization tool for atmospheric data, the user provides information about the data for which a tool is to be constructed. The facility for entering this information is called the Object Editor, shown in Figure 1. This mouse-driven tool is used to enter the size, orientation, and format of the data, grid geometry or unstructured connectivity information, and file locations. At this point, the user gives each field in the data set a descriptive name and identifies it as vector or scalar. The user then gives the data description a name and saves it to a file. The Object Editor uses this information to create a visualization tool whose back-end processing is performed on any specified host.

1.2 User Filters

Often, large scientific computation data sets are limited in the number of dependent variables that are available from the base computation. For example, a large CFD computation may contain only the basic velocity field vectors in its output. If during post process examination, vorticity

information is required, additional large-scale processing usually is needed. We have developed a new type of data structure that allows very flexible access to large data sets. By coupling this data structure with the equation editor we can offer the user the ability to generate interactive filters for these large data sets. This reduces the need for large-scale recomputation [Peskin et al., 1992].

User-added filters are specified as vector or scalar operations on existing fields. Filters are entered as mathematical expressions which are parsed and passed to a symbolic solver that understands the structure of the data. New fields are computed from the user-added filters based on interpolations of the existing fields across the domain. Once computed, the tool treats a field generated by a user-added filter as it would the original data. It may be visualized graphically, or used as the basis of further calculation. The numeric values of a field generated by a user filter may be examined within the tool.

1.3 Grid Issues

SCENE supports meshed data in structured and unstructured formats. For structured grids, general

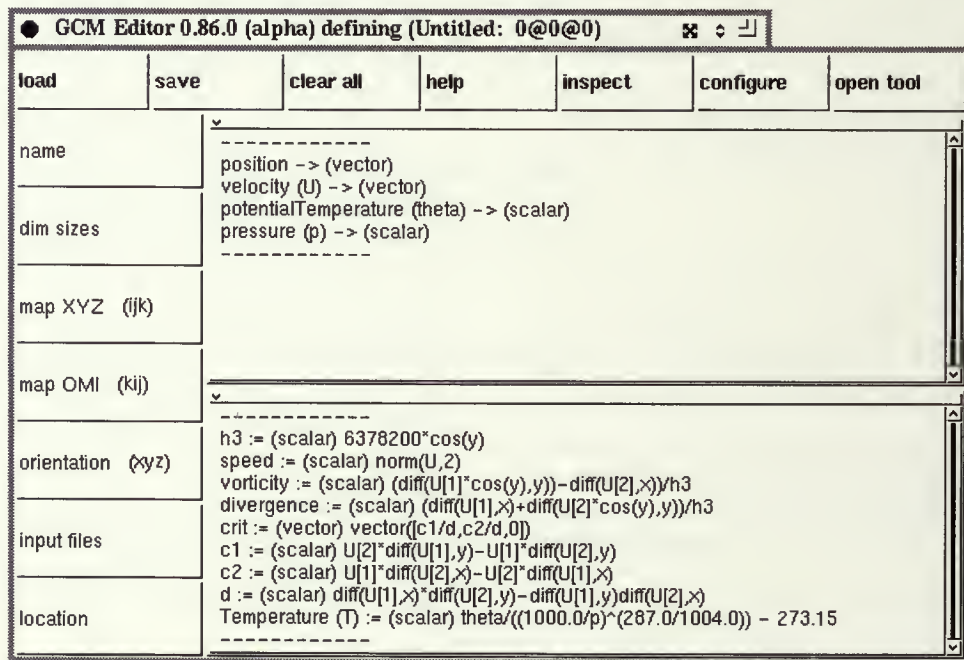


Figure 1. Object Editor used to create a SCENE visualization tool for a General Circulation Model data set. The boxes on the top and left indicate the structure of the data files and how they are to be displayed. The fields to be visualized from the data files are shown in the center. Below these are equations entered as user-added filters. The syntax is that of the MAPLE symbolic computing package.

curvilinear coordinates are supported. At every point in the structured grid, the directional derivatives of the mesh are calculated to form a local Jacobian of the coordinate transformation. The grid metrics are used to specify geometrically proper vector operators, such as grad, div, and curl [Greenberg, 1978]. Proper formation of the grid metrics is a prerequisite for integrating particle pathlines.

For general curvilinear coordinates, the partial derivatives are computed at every point in the grid using five-point differencing. These define a 3×3 matrix, which is the inverse of the Jacobian of the coordinate system. The inversion of this matrix at every point is not particularly expensive, because the matrix is small and the routine only needs to be performed once for a given grid.

For unstructured meshes, connectivity information about the mesh must be supplied. The basic elements of the two-dimensional unstructured mesh are triangles, each of which must contain information about its vertices and its neighbors (triangles that border its sides). Higher order computations are possible only on meshes that contain still more connectivity information [Froncioni and Peskin, 1993].

1.4 Visualization Options

From the Object Editor, a mouse click will launch a visualization tool for the data. All plotting options are available on fields generated from user filters as well as the stored data. The functionality is the same for gridded and unstructured meshes. The tool has four menus, which set visualization options. Menu one controls file input and output. As soon as the tool is created, the user loads the specified data in the Object Editor. If the user filter is used to calculate additional fields, these may be saved to files. Menu two sets the plotting option. Available graphical views include contour plots, vector plots, particle tracking, wireframe drawing, and surface plots. Menu three controls subview options for row plots, column plots, zooms, and point probes. Menu four is used to set the field that is to be visualized. Figure 2 shows several of the available options.

In a SCENE visualization, the graphical view is always linked directly to the underlying data. The user can probe the numerical data by clicking the

mouse on the view. A window will pop up, which displays the numeric value and type (vector or scalar) of every field at that point. A subview may be created of any portion of a graph by dragging the mouse pointer over an area of the view. The subview is not merely a magnification of the original drawing, but rather a redraw that might indicate features too small to be resolved in the original view. Zoomed views have the same functionality as the original graphs, including the data probe. The user also has the option to zoom again in the subview and to zoom in on any number of different regions from the original view.

2. PARTICLE TRACKING

2.1 Trajectories in Steady and Unsteady Flow

SCENE has the capability to trace the field lines of any vector in the data set. The field lines of the velocity vector describe the trajectory of a massless particle for steady flow, which defines a streamline. For a time-dependent data set, the user has the option to trace true pathlines in the unsteady flow or streamlines at an instant (Figure 3). Pathlines show the time history of single particles as they are advected with the flow. If the flow is unsteady, particle trajectories may cross themselves, while streamlines never cross. In Figure 4, the red pathline and the blue streamlines are traced from the same region of the western Pacific Ocean. The trajectory of the pathline is similar to one of the streamlines, but the unsteady tracer moves erratically. In steady flow, the streamlines and pathlines coincide.

2.2 Trajectory Options

The particle-tracing option (path tool) has functions that provide flexibility to the user. For example, the starting coordinates of a pathline may be specified from a pop up menu, which allows highly accurate specification of the location. The user may instead start pathlines at a location chosen with the mouse. A third option is to specify starting locations in a file, which is useful for tracing a set of pathlines repeatedly. If starting locations are read from a file, the color of each line may be specified as well.




Figure 3. Streamlines at the 0.25-kg/m^3 density altitude of the GISS-GCM for a typical model year. Velocity field lines are traced from points with roughly 10 degrees of latitudinal separation. The streamlines demonstrate that the flow is predominantly zonal in the middle latitudes, while the tropics are more complex and contain many recirculation zones.

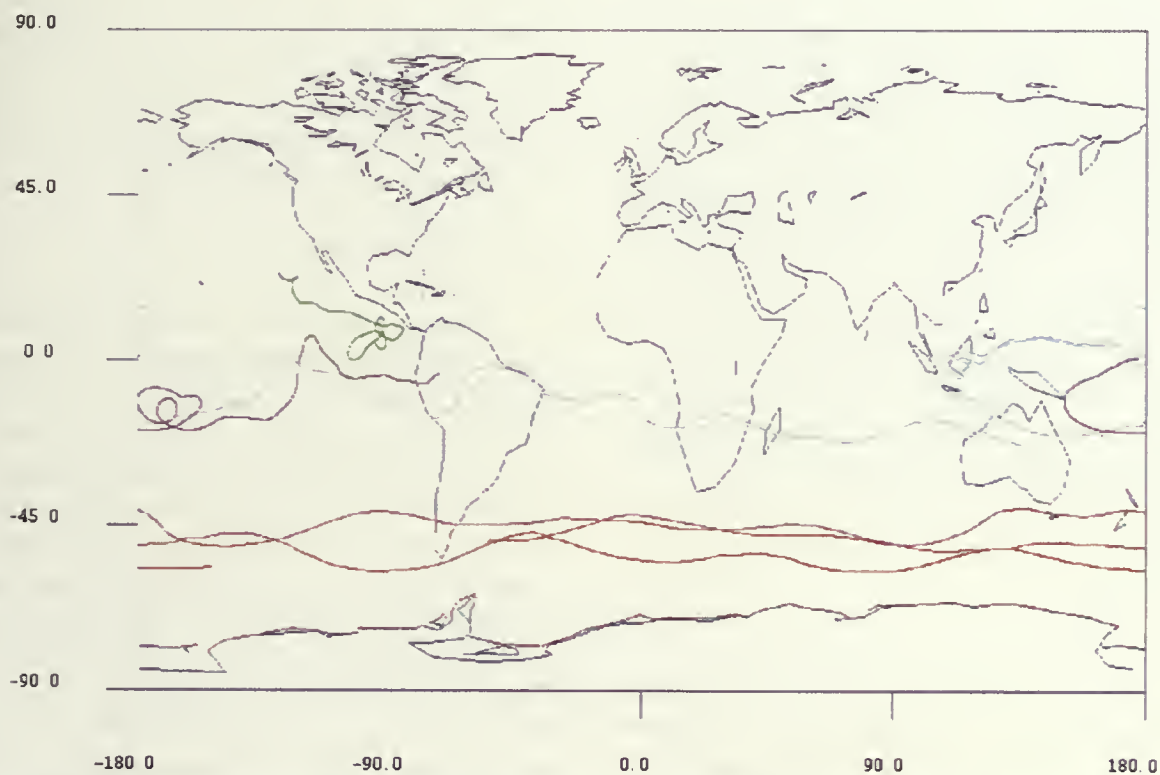


Figure 4. Pathlines (red, magenta, and green) and streamlines (cyan) for the 0.25-kg/m^3 density altitude of the GISS-GCM. These tracers were released interactively by the user to locate interesting regions of the flow in both the transient and instantaneous domains.

The time step for trajectory integration is adjustable, and both positive and negative step sizes are valid. Pathlines are traced by integrating in time from any starting point. The usual method of integration is fourth-order Runge Kutta, but lower order integration methods have been considered for situations where computing resources are limited. To use the Runge Kutta method, four interpolations into the grid are required for each time step, and each interpolation is numerically expensive. As an alternative, a second-order, modified Taylor method suggested by Westerveld (1989) has been incorporated into the path tool for general curvilinear coordinates. For this method, derivatives of the velocity field are precomputed and stored at each grid point.

3. VISUALIZATION OF GENERAL CIRCULATION MODELS

As part of a study comparing numerically calculated particle trajectories with balloon observations from the TWERLE experiment [Er-El and Peskin, 1981], a SCENE visualization environment has been created for data generated by the General Circulation Model at the Goddard Institute for Space Studies (GISS-GCM). The purpose of the GCM tool is to explore the transient flow field and analyze the motion of many particle trajectories. The model uses a grid of 36 lines of longitude by 24 lines of latitude with nine atmospheric levels. A finer resolution is also available with 72 longitudes and 46 latitudes. The grid covers the entire surface of the Earth, and the nine vertical layers extend from the planetary surface to well above the jet stream. The Earth is gridded in spherical coordinates with a constant radius. The available data contain 12 monthly averages to simulate a typical calendar year. Another set contains 31 daily averages for January 1979. The nine vertical coordinates are set at constant levels of pressure difference.

3.1 Data Representation

Three coordinates are required to specify a point on the surface of the Earth in Cartesian coordinates (such as length, width, and height from the Earth's center). If spherical coordinates are

used instead, only two coordinates (longitude and latitude) are necessary. This is a substantial advantage, but it creates a problem of grid degeneracy at the poles. The existence of poles prevents SCENE from using the techniques developed for general curvilinear coordinates. For this reason, support for spherical coordinates is a separate option that is built into the GCM visualization tool. This option provides the user with the ability to perform vector operations that are even more accurate than those for curvilinear coordinates because the grid metrics are described analytically. To account for the non-affine mapping of the Earth onto a rectangle, the interpolation is weighted by projected area. This correctly accounts for the latitudinal variation of a cell's area.

In addition to the basic advection equations, GCM integrates the energy sources to the flow whose effects may be highly localized, while the resolution of the model is quite coarse. For this reason, first-order interpolation is used wherever possible. Higher order interpolation, which sacrifices locality for higher nominal accuracy, may produce undesired smoothness in the results. For example, GCM's differing representations of sea and land surfaces generate a steep gradient at the coastlines, especially at the mountainous west coasts of the American continents. These gradients persist into the higher levels of the atmosphere and the effects of these gradients are diminished when higher order interpolation is used.

3.2 Data Reduction

The volume of data taken from GCM was quite large—46 latitudes x 72 longitudes x 9 levels x 31 days. The two vector field's position and velocity were stored, along with two thermodynamic scalars (pressure and potential temperature). In single precision, this is 29 megabytes of data. While SCENE is able to handle large data sets, only one plane of constant density altitude is of interest for comparison with the TWERLE experiment. Also, the GCM data are generated on surfaces of constant relative pressure difference (sigma levels), so the data are not directly comparable with the experimental results. To trace Lagrangian particles that model the balloons accurately, isosurfaces of density were formed by interpolating between the

pressure surfaces. One constant density surface was formed for each day, thus reducing the dimensionality of the problem to three: latitude, longitude, and time. The storage is therefore $(2 \times 3\text{d vectors} + 2 \text{ scalars}) \times 46 \times 72 \times 31 \times (4 \text{ bytes per float}) = 3,285,504 \text{ bytes}$, which is managed by SCENE running on a single workstation.

4. QUALITATIVE DESCRIPTIONS OF ATMOSPHERIC FLOW

The global atmospheric flow field contains large-scale structures that are analyzed with the GCM tool. The primary means of describing the instantaneous flow field is streamline tracing. Figure 3 shows a streamline pattern created with the path tool at one time step. The expected pattern of smooth flow at the middle latitudes and recirculation in the tropics is apparent. Terminating streamlines indicate flow sources and sinks in the two-dimensional flow field. A streamline that spirals into a stagnation point may indicate a true sink in the compressible flow field. However, the flow field of Figure 4 is at a constant density, so the streamlines terminate where the flow is predominantly vertical.

4.1 User-Added Filters on GCM Data

The vertical component of the GCM velocity field is a diagnostic variable that was not included in the data set. The GCM tool was therefore used to generate information about the flow in this direction. For incompressible flow, conservation of mass indicates that the vertical flow is found by integrating the continuity equation over the thickness of the atmospheric layer. The term inside the integral is the divergence of the two-dimensional flow field, the equation for which is entered as a user filter in Figure 1. The full solution for the vertical velocity requires a constant of integration at every grid cell, but the information from GCM required to calculate this constant was not included in the available data set.

4.2 Transient Flow Description

The transient properties of the flow field are less easily analyzed, partly because of the higher

dimensionality of the data. While it is not difficult to create a contour plot with time as one coordinate direction, the interpretation of such a picture is not always apparent. Also, streamlines are informative, but they do not give a complete description of a transient flow. The primary means of exploring time-dependent properties of the flow is with pathlines. The interactive nature of the tool is important when tracing pathlines; the static picture of Figure 4 is not ideal for portraying a time-varying simulation.

Interesting transient behavior is shown by the green pathline west of Central America. A massless particle released at the first time step recirculates in a very small region for most of the simulation before moving off to the northwest at the end of the month. The figure demonstrates a primary distinction between streamlines and pathlines—the streamlines are isolines of a vector field and cannot cross themselves or each other in the two-dimensional field. Pathlines have no such restrictions, as each line is the projection of two-dimensional plus time trajectory onto the plane. Streamlines are contrasted with pathlines in the figure by starting light blue streamlines and one pink pathline in the same region of the equatorial Pacific. The pathline moves erratically but follows the general flow pattern indicated by one streamline. The other streamline falls into a sink in Indonesia, which exists in the July 1 mean flow field. The red pathline indicates that the flow in the middle latitudes is smooth and predominantly longitudinal for the entire month.

4.3 Critical Point Analysis

The streamlines shown in Figure 3 were traced from starting positions separated by five degrees of latitude. The starting points for these were chosen by the user arbitrarily, and the resulting picture misses important details. In the tropics, more streamlines are required to resolve smaller scale structures, while many streamlines in the smooth flow of the mid-latitudes are redundant. The identification of flow topology with greater precision is possible by examining the critical points of the flow, as demonstrated by Dickinson (1991) and Froncioni and Peskin (1993). The techniques of critical point analysis are applied in

the GCM tool only as a useful example. A complete implementation requires the construction of a connectivity diagram that identifies all critical points in the data and the streamlines that connect them.

The critical point equation, applied as a user filter in Figure 1, defines a vector field that is meaningful only near a saddle point. In this region, the field points to the saddle so that a field tracer started in this region will move directly to the saddle point. A close look at the streamline pattern of Figure 3 indicates the existence of a saddle point in the tropical Atlantic Ocean. Using the mouse, a location off the African coast is selected, and a line of the newly calculated critical field is traced, which is shown in red in Figure 5. The field line terminates at the saddle point to the west. To verify that this is the proper location of the saddle point, two streamlines are shown in green. The lines are initially close to each other and are well correlated as they travel eastward over South America. However, the streamlines diverge immediately as they pass on opposite sides of the saddle point.

5. CONCLUSION

We have presented the SCENE interactive data exploration environment and demonstrated its applicability for the manipulation of large multidimensional data sets. The use of an object-oriented program structure makes it possible to create extensible, customized visualization tools. This flexibility is achieved without any programming required by the user.

This is demonstrated with a visualization tool for the exploration of data from numerical general circulation models. The simulated general circulation can be visualized in several ways. At a specified time, the instantaneous flow field can be explored with streamlines in combination with critical point analysis. Derived fields can be generated and displayed interactively with user-added data filters. The graphic displays remain linked with the underlying data, which enables examination of the data by clicking on the picture. The time-dependent flow field was investigated qualitatively with pathlines.

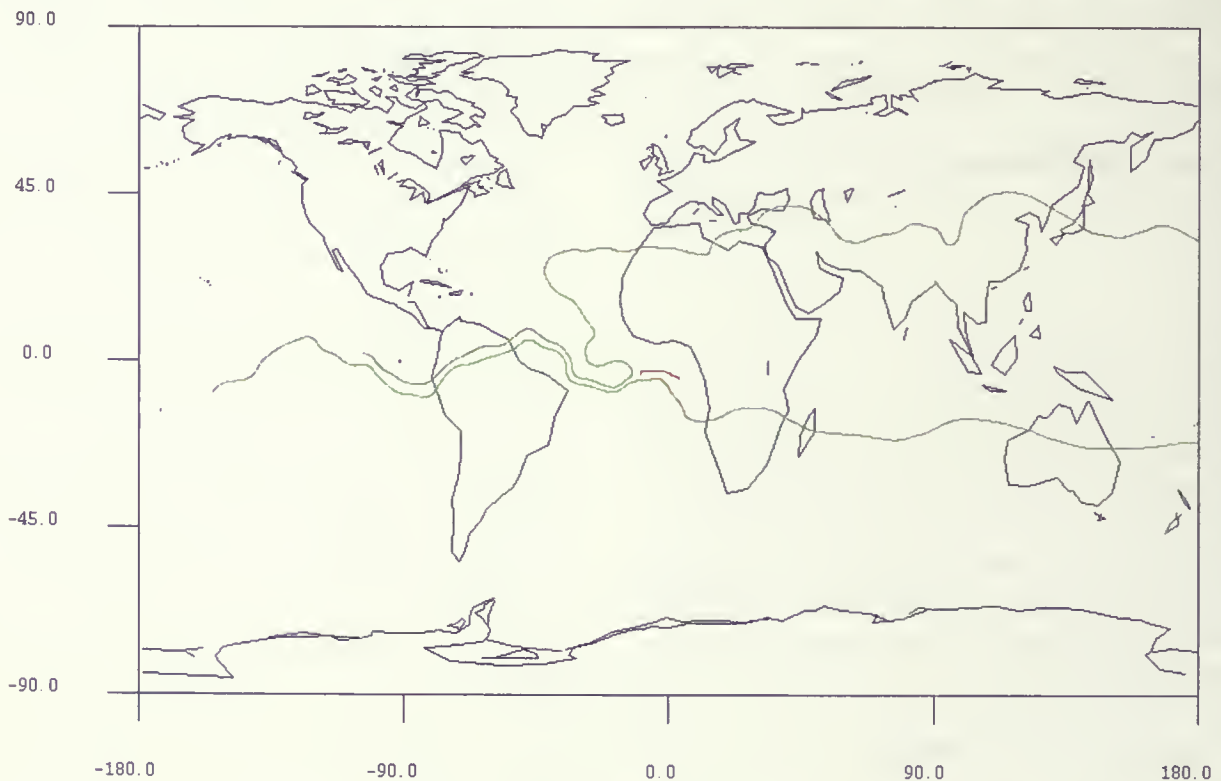


Figure 5. Saddle point locator (red) and streamlines (green) for the 0.25-kg/m³ density altitude of the GISS-GCM, for a typical daily mean flow field. Two streamlines, initially similar, diverge as they pass on either side of the saddle.

SCENE is available on the Internet via anonymous FTP to cesl.rutgers.edu. A multimedia demonstration will soon be accessible via WWW. For further information or details about obtaining and configuring SCENE, contact rosen@caip.rutgers.edu.

REFERENCES

- Danielsen, E.F., Review of Trajectory Methods, *Advances in Geophysics*, A, 18, 73-94, 1974.
- Dickinson, R.R., Interactive Analysis of the Topology of 4D Vector Fields, *IBM J. Research and Development*, 35, 1, 59-66, 1991.
- Er-El, J. and R.L. Peskin, Relative Diffusion of Constant-Level Balloons in the Southern Hemisphere, *J. Atmos. Sciences*, 38, 2264-2274, October 1981.
- Froncioni, A.M. and R.L. Peskin, Qualitative Flow Descriptors for Unstructured Triangular Grids, *Third International Conference on Expert Systems for Numerical Computing*, in press, 15 pp., West Lafayette, IN, 1993.
- Greenberg, M.D., *Foundations of Applied Mathematics*, 175-177, Prentice-Hall, Englewood Cliffs, NJ, 1978.
- Hung, C-M. and P.G. Buning, Simulation of Blunt-Fin Induced Shock Wave and Turbulent Boundary Layer Interaction, *AIAA 22nd Aerospace Sciences Meeting*, 19 pp., Reno, NV, January 1984.
- Liu, M., *Numerical Study of Nonlinear Dynamics of Particles and Chaotic Mixing in Two Dimensional Periodic Driven Cavity Flows*, Ph.D. thesis, Rutgers University, New Brunswick, NJ, 1992.
- Peskin, R.L., S.S. Walther, and T. Boubez, Computational Steering in a Distributed Computer Based User Interface System, *Artificial Intelligence, Expert Systems and Symbolic Computing*, 104-113, edited by E.N. Houstis and J.R. Rice, Elsevier Science Publishers B.V. (North-Holland), 1992.
- Pierrehumbert, R.T. and H. Yang, Global Chaotic Mixing on Isentropic Surfaces, *J. Atmos. Sciences*, 50, 15, 2462-2480, 1993.
- Walther, S.S. and R.L. Peskin, Parallelized Interactive Data Management of Large, Scientific Data Sets, *Proceedings of Sixth SIAM Conference on Parallel Processing for Scientific Computing*, 4 pp., March 1993.
- Westerveld, W.B., Plotting of Closed Field Lines With an Application to Coherently Excited Atoms, *Computers in Physics*, 74-77, 1989.
- Zinn, H.P., *An Efficient Boundary Layer Model for GCMs: Its Design, Validation, and Implementation Into the GISS-GCM*, Ph.D. thesis, Rutgers University, New Brunswick, NJ, 1993.

A Data Reduction, Management, and Analysis System for a 10-Terabyte Data Set

R. DeMAJISTRE

Computational Physics, Inc., Fairfax, VA

L. SUTHER

The Johns Hopkins Applied Physics Laboratory, Laurel, MD

Within 12 months, a 5-year space-based research investigation with an estimated daily data volume of 10 to 15 gigabytes will be launched. Our instrument/analysis team will analyze 2 to 8 gigabytes per day from this mission. Most of these data will be spatial and multispectral, collected from nine sensors covering the UV/Visible/NIR spectrum. The volume and diversity of these data and the nature of its analysis require a very robust reduction and management system. This paper is a summary of the system's requirements and a high-level description of a solution. The paper is intended as a case study of the problems and potential solutions faced by the new generation of Earth observation data support systems.

1. INTRODUCTION

During the next decade, a new generation of Earth-observing platforms will be acquiring large and diverse data sets. To analyze these data effectively, new data management and reduction techniques must be developed. We present here a case study of the data reduction, management, and analysis system currently under development for the Ultraviolet-Visible Imaging and Spectrographic Imaging (UVISI) instrument aboard the Midcourse Space Experiment (MSX). This instrument will be used for Earth and celestial observations and has an estimated accumulated data volume of 10 terabytes over the 5-year mission. The following sections include a description of the platform and the instrument to be supported and the data set itself, a summary of the data analysis requirements, and a discussion on the architecture of the system under development.

2. THE DATA SOURCE

The MSX spacecraft will be placed in a polar orbit at an approximate altitude of 900 km. The spacecraft carries several optical remote sensing instruments, including the UVISI instrument, which is our payload of interest. Aside from small line-of-sight modifications affected by scan mirrors, all of the optical sensor pointing is determined by the attitude of the spacecraft. All of the optical instruments share a nominal optical axis, which can be pointed in an arbitrary manner.

Data from the satellite are downlinked at a single ground site at a rate of 25 megabits per second. The total data volume is therefore limited by contact time with the ground station. It is estimated that the average daily data output will be

10 to 15 gigabytes. Onboard tape recorders can record data at both 25 and 5 megabits per second. In the 5-megabit mode, data can be collected for more than 5 hours per day. The daily data rate of UVISI, the payload of interest, is determined by mission priorities. UVISI shares the downlink and other resources with an infrared instrument with a nominal 18-month lifetime. During this period, the infrared instrument will have priority on the link, and the UVISI data stream is estimated at 2 to 5 gigabytes/day. For the remaining 42 months of the 60-month mission, the data rate for UVISI is expected to regularly exceed 8 gigabytes/day.

UVISI is composed of nine sensors, consisting of four spatial imagers and five spectrographic imagers (SPIMs) [Carbary, 1992]. Each of these sensors has its own optics and intensified charge coupled device (ICCD) detector. The intrascene dynamic range for these sensors is 12 bits (4096). The interscene dynamic range is determined by the intensifier gain and is approximately 106. The framing rate for all of the sensors is either 2 or 4 hertz. The spatial imagers provide a variety of angular coverages and spectral sensitivities. Filter wheels can be used to select spectral bands. The spatial imagers do not employ scan mirrors, and are thus completely dependent on spacecraft attitude for pointing. The SPIMs employ dispersive optics to create spectral images, which contain both spectral and spatial information. Each SPIM can acquire up to 40 spectra of 272 channels simultaneously. The field of view of the SPIMs is much more narrow than that of the imagers. However, the field of regard is increased by the use of scan mirrors. The SPIMs provide continuous spectral coverage from the far ultraviolet (FUV) to the near infrared (NIR) (100 to 900 nanometers). The large dynamic range, spectral coverage, and variety of instrument modes provide the investigators with the means to explore a great number of different phenomena.

3. THE DATA SET

The large data rate and the flexibility of both satellite pointing and observation modes make UVISI a very challenging data set to manage and analyze. Because of the size and complexity of the data, recalibration and data location are expensive

tasks. For this reason, several different types and levels of data products are planned.

The primary archive of instrument data will be composed of minimally processed raw data (termed level 1a data). The minimal processing resolves overlaps, separates the data by instrument, and places the data in the proper time order. The data are not stored in a calibrated form for two reasons. The primary reason is the expense of recalibrating large portions of the data set. With the present scheme, only the fraction of the data that is of interest to the investigators needs to be calibrated. Furthermore, the process of calibration generally expands the size of the original raw data set, by converting integer values to the floating point and specifying the uncertainties of the calibrated data. It has been estimated that the UVISI data will expand by a factor of 8 when calibrated. The storage of calibrated data is therefore much more expensive.

To be useful, the data must be calibrated. For UVISI, a standard set of calibration routines and data is supplied to our facility by the UVISI instrument team [MSX, 1991]. All calibration changes are then reflected in the software or calibration data. Thus, only the most important data need to be recalibrated. The output from the calibration software is termed level 2 data. These data contain radiometrically calibrated images but do not contain processed pointing information. In addition to the calibrated data, level 2 files contain indicators of instrument configuration and data health.

Data location also is a serious challenge to the analysis of the UVISI data. To facilitate the location of useful data, summary files are produced for each data set. These summary files contain sensor configuration and data quality information, as well as calibrated radiometric summaries of the data taken. For the spatial imagers, the radiometric summaries include mean pixel values and overruns for different portions of the focal plane. For the SPIMs, the total intensity measured within a set of given spectral ranges is reported. These radiometric values are much smaller than the images themselves and can be used for rapid searches. All of the summary products are produced "upstream" of our facility, and the details of their creation will not be described here.

Knowledge of the radiometrics is of limited utility without sensor pointing data. Corrected

ephemeris and attitude information for the MSX spacecraft is supplied by a processing center dedicated to this task. Our facility is provided with the spacecraft position and the attitude (in quaternion form) of the nominal optical axis of all the instruments and with sensor offsets from this axis. We can thus determine the attitude of a fiducial point for each of the UVISI sensors. It is our task to convert this set of ephemeris attitude data into useful pointing quantities (such as spacecraft latitude, tangent point altitude, etc.) for either the center of each imager or spatial pixel. It is also our task to merge the pointing data with the radiometrically calibrated data.

The products outlined above will provide most of the information necessary for our data reduction, management, and analysis. Low rate information, such as planning data and spacecraft housekeeping, will also be used but will not be central to our processing.

4. ANALYSIS REQUIREMENTS

The flexibility of UVISI allows a large number of phenomena to be studied. The investigators at our facility have defined 15 different experiment plans, ranging from auroral studies to lightning measurements. These plans define dedicated observations by UVISI. In addition to these dedicated observations, other investigation teams will use UVISI for celestial observation, for contamination studies, and as a complement to the infrared observations. Data taken during these other observations also may be of use to our team. Furthermore, several analysis tasks have already been defined that combine data from planned observations. These are essentially plans to re-use the data taken during the dedicated observations for alternate purposes. It is expected that the definition of these alternate analysis plans will continue throughout the mission, thus requiring us to be able to identify data that are useful to such tasks. We have thus separated the UVISI data analysis into three types of processing tasks:

- Planned experiment processing
- Planned re-use processing
- Historical re-use processing

The principle use of this classification is data identification. Data from the planned experiments are

selected on the basis of the mission's operations schedule. The planned and historical re-use data are identified by instrument pointing, configuration, and other metadata describing the observations. In the case of planned re-use data, the data are selected from the incoming data stream. The historical re-use data involve extracting data from the UVISI archive.

Once the data for a particular analysis task have been identified, they must be reduced or otherwise processed according to the requirements of that task. Most of the tasks have different processing requirements. After the data have been prepared for analysis, they must become accessible to the investigators in an interactive mode. Data reduction and analysis thus fall naturally into two categories: (1) pipeline data production and (2) interactive analysis. The investigators have been encouraged to place any analysis processes that can run in an unattended, automated mode in the pipeline for that task. This will allow the investigators to concentrate on the less stable elements of their analysis. It is expected that as the mission progresses and the analysis algorithms and the instrument data become better understood, the processes that had been interactive will migrate to pipeline processing. After the data have been through the pipeline, it will be available for interactive analysis. A schematic of the data flow is illustrated in Figure 1.

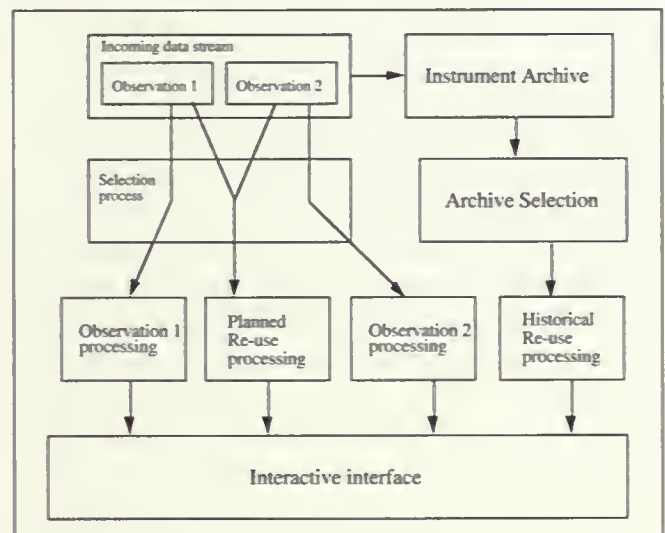


Figure 1. Data separation and processing.

5. ARCHITECTURE

The high-level architecture of our system is presented in Figure 2. The four basic subsystems

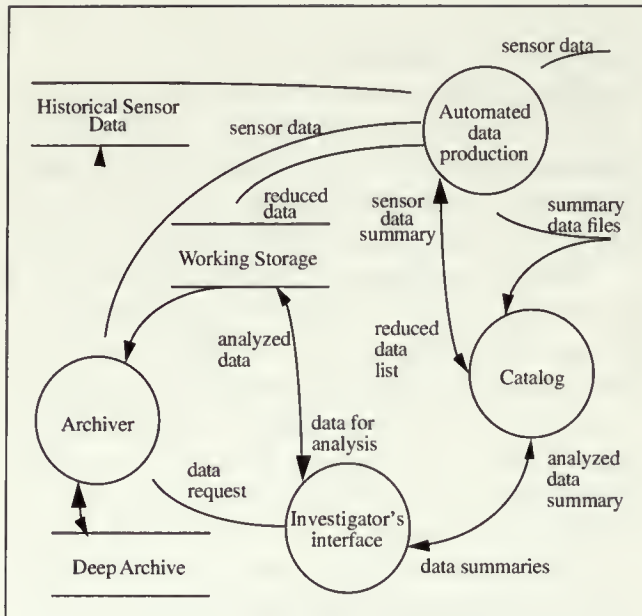


Figure 2. Data flow for major system components.

are: automated data production, catalog, investigator's interface, and the archiver. When data are received, all of the incoming data are made available to the automated data production subsystem. The summary files and attitude information are passed in parallel to the catalog subsystem. The catalog subsystem builds a unified description of the sensor data from the files that it receives and passes this information back to automated data production. Based on this unified summary, the data production system decides how to process the sensor data. The data are processed, and the output of the processing is placed in working storage. A summary of these products is sent to the catalog subsystem. The data have a limited lifetime in working storage and will be placed into the deep archive when it is not being used. The archiver subsystem manages data transfers between working storage and the deep archive. The archiver may also place data in the historical sensor data area, which can be used as an alternate input to automated data production. The investigator's interface is an end-user graphical user interface to the data and a set of analysis tools. It inputs data directly from working storage; historical analysis can be initiated by searching the catalog summaries for the data of interest and requesting their promotion to working storage from the archiver. Any analysis products that are deemed suitable for permanent storage can be placed into working storage for subsequent

archival (a summary of these data must be placed in the catalog for future access).

The automated data production subsystem is based on the individual analysis tasks. Each task is associated with a processing "box." These boxes are constructed based on the requirements of their respective tasks. During the construction of the boxes, we have built a library of common data manipulation tools, which are also shared with the investigator's interface. Each box is designed to run independently and in an unattended mode. The subsystem also contains a "box integration" element, which selects data for the boxes, schedules their execution, and organizes their output in working storage. In effect, the box integration element acts as a shell that determines how the boxes run. The operational scenario of the automated data production subsystem is as follows. Upon receipt of new data, box integration queries the catalog to determine which boxes, if any, should be applied to the new data. The sensor data are then sent to working storage for archival. Box integration then organizes the data appropriate to each selected box and then executes that box. The boxes reduce the data, reporting their progress in a status log, managed by box integration. Upon completion of box processing, the output products are then organized appropriately in working storage. A summary of these products is also produced by box integration and is posted to the catalog. Thus the data have been reduced, described, and prepared for archive.

The subsystem catalog contains a description of all the products that are available at our facility. The relationships among the products are also described by the catalog, so that a unified description of the facility data holdings is always available. Summaries of the sensor data and derived products are constructed from the summary files and the attitude information. Note that the attitude information is converted into useful pointing quantities as it is loaded. The use of the summaries not only provides an excellent description of the data holdings, but also concisely describes the mission itself. Thus, the catalog can provide a history of the mission as well as assist data location. We have separated the catalog into two parts: the staging catalog and the historical catalog. The staging catalog contains only the newest data (these data

have a lifetime of about 7 days). The historical catalog contains data summaries from the entire mission. The staging catalog is used by the automated data production subsystem when it is processing new data. It also serves as a quarantine area and transaction buffer to protect the historical catalog from badly formatted input data and to allow for more efficient population of the larger catalog. The historical catalog will allow the users to survey the entire data holdings without having to access the data in the deep archive. The estimated size of the deep archive is 10 to 12 terabytes, whereas the catalog will only require about 10 to 20 gigabytes, and can remain online throughout the mission. The historical catalog is used by the automated data production system for selection of historical sensor data for re-use processing. It can also be used by the investigators to determine whether the current data holdings can usefully support a perspective analysis task. The historical catalog adds much value to the holdings of our facility, in that a small amount of useful data can be located and extracted from the large archive. Without the catalog, much of the data would rot on the shelf.

The investigator's interface subsystem is a graphical user interface that provides access to all levels of the data, as well as the means to analyze them. Various data analysis tools, developed in concert with the box processing, can be used interactively or be assembled into batch processes. The library of tools can be extended by the users and may include specialized tools necessary for particular analysis tasks. The interface includes several functions for locating, reading, and writing data in working storage. Final products constructed during interactive analysis can be registered in the catalog and stored. Intermediate and working products can be placed in a user's own work space. The interface will include access to several graphical display and analysis packages, such as IDL, AVS, and Khoros. This will allow users to analyze their data with familiar tools without having to worry about writing tedious input/output routines.

The archiver subsystem moves data to and from the deep archive. The bulk of the facility's data holdings will be stored in the deep archive. The current primary storage medium for the deep archive is tape storage. The archiver is a combina-

tion of the automated process of tracking data files and their location on physical media and the manual process of walking between the shelf and the tape drive. The archiver process maintains a data base of the contents of the deep archive and the working storage area. It also is responsible for migrating data to and from working storage. Data will reside in working storage only as long as they are actively in use. After a given time of inactivity, if the data have not already been placed in the deep archive, they are archived and deleted from working storage. Data that already reside in the deep archive are simply deleted from working storage. Reliable performance of the archiver is crucial to the success of the project, so every effort has been made to keep its design as simple as possible.

6. SUMMARY AND DISCUSSION

We have presented a system architecture for the reduction, analysis, and management of large volumes of Earth observation data. This architecture is sufficiently flexible to allow for changes in instrument operation and calibration, as well as changes or additions to the reduction and analysis algorithms applied to the data. Our primary goal in the construction of the UVISI processing system is to allow the investigators to concentrate on the content of the data rather than the details of the data processing system.

By far, the most challenging elements of manipulating a data set of this size (approximately 10 terabytes) are data location and flexible processing. Unless robust data location methods are implemented, the task of finding data useful for a particular analysis becomes impractical. Flexibility of processing is vital to research investigations such as these. More often than not, data processing systems have to be significantly modified because of unforeseen circumstances after data acquisition had begun. The utility of the data reduction, management, and analysis system cannot be over-emphasized. A properly implemented system allows the investigators to maximize the utility of the data set and to apply previously acquired data to new problems. Without this capability, newly formulated problems will require new missions, which in turn will acquire large volumes of only marginally useful data.

REFERENCES

Carbary, J., UVISI Instrument: A Tutorial, APL S1G-R92-04, June 1992.

Midcourse Space Experiment (MSX) Preliminary Data Management Plan, MSX Program Office, DM-070101-p02.01-910715, September 1991.

Geophysical Data Analysis and Visualization Using the Grid Analysis and Display System

BRIAN E. DOTY and JAMES L. KINTER, III
*Center for Ocean-Land-Atmosphere Studies, IGES,
 Calverton, MD*

Several problems posed by the rapidly growing volume of geophysical data are described, and a selected set of existing solutions to these problems is outlined. A recently developed desktop software tool called the Grid Analysis and Display System (GrADS) is presented. The GrADS' user interface is a natural extension of the standard procedures scientists apply to their geophysical data analysis problems. The basic GrADS operations have defaults that naturally map to data analysis actions, and there is a programmable interface for customizing data access and manipulation. The fundamental concept of the GrADS' dimension environment, which defines both the space in which the geophysical data reside and the "slice" of data which is being analyzed at a given time, is expressed. The GrADS' data storage and access model is described. An argument is made in favor of describable data formats rather than standard data formats. The manner in which GrADS users may perform operations on their data and display the results is also described. It is argued that two-dimensional graphics provides a powerful quantitative data analysis tool whose value is underestimated in the current development environment which emphasizes three-dimensional structure modeling.

1. INTRODUCTION

The earth scientist today faces an exponentially increasing volume of observational and model output data. The fraction of that data that can be sampled, organized, archived, and intelligently analyzed grows constantly smaller as the improvement of computer hardware and software fails to keep pace with the growing volume of data.

While it is true that computer hardware development has advanced rapidly, there are a number of problems with the software that has been and is being developed for the high-end workstation platforms. A large number of data analysis and visualization programs are available in the form of commercial off-the-shelf software (COTS). Many of these programs hold great promise for solving the data volume problem, but they have a number of disadvantages, which make them inaccessible to a large portion of the earth science community. Many COTS products are too expensive or too difficult to learn for many scientists who have limited funding or time budgets. Several COTS programs support only their own internal data format or a limited number of "standard" or "common" data formats so that a given scientist's data may not be handled by these programs. Most earth scientists have found that 80386- or 80486-based computers are adequate for their work and are consistent with their equipment budgets, but COTS developed for RISC/UNIX computers or for special purpose hardware (such as three-dimensional graphics acceleration boards), will not run on these machines. Also, in our opinion, there has been an overemphasis on three-dimensional structure modeling and object rendering, which, admittedly, provides a useful tool for intuitive or conceptual exploratory data analysis but is seriously lacking in

the kind of quantitative data analysis scientists require on their desktops that is provided by two-dimensional graphics. Finally, it is our opinion that COTS products do not offer all that is needed to all earth scientists because they typically have a steep learning curve and require sophisticated programming (such as the construction of modules and module networks in the case of data flow environment software).

Existing visualization software is frequently less useful to earth scientists than it could be because it is designed and written with a singular emphasis on computer graphics or an "ergonomic" user interface that purports to be intuitive by allowing users to navigate through nested menus or linked lists of point-and-click mouse button events. The result is that insufficient design attention is paid to the scientific problems at hand. In particular, existing software does not adequately link together the processes of data analysis and visualization. In cases where analysis is emphasized, scientists must export their results to other platforms or software packages to depict them graphically. When visualization is emphasized in a given program, there is no capability for quantitative analysis or comparison among results, and the linkage between the data and the geophysical context (map background, earth-registered coordinates, etc.) is missing. Although we make these assertions without proof, there are some articles and scientist surveys that support them (such as Doswell, 1992; Botts, 1993, personal communication).

Our experience is that earth scientists are only briefly enamored of the graphical user interface (GUI) to a particular data analysis and visualization program when they first begin using the program. After becoming familiar with the program, they find that the multitude of menus, button clicks, and dialog boxes slow down their productivity. Many earth scientists desire graphical user interfaces to their own data—that is, they want to have a graphical means of accessing, manipulating, and displaying their own data that is intuitively familiar and consistent with the way in which they normally approach their data. Thus, a program GUI, which provides users a graphical means of accessing all the myriad features of a visualization program, often gets in the way, but a data GUI, which pro-

vides a simple means for users to work with their own data, is highly desirable.

Finally, many earth scientists have taken exception to existing programs because they are insufficiently extensible or customizable and because they divide the tasks of data analysis and visualization rather than linking them more closely together. Earth scientists want to have the capability to add functionality, add their own analysis techniques, and customize existing programs to their data. They have traditionally done this with FORTRAN and the programmable (batch) graphics software packages (such as NCAR Graphics). Earth scientists also have integrated data analysis with visualization in the past by writing FORTRAN (batch) programs that include both calculations involving basic data sets and calls to graphics libraries to display the results of those calculations. They want to be able to achieve the same level of integration interactively so that the data analysis process can be streamlined and made more intuitive.

To provide a means for conducting geophysical data analysis interactively, we have designed and implemented the Grid Analysis and Display System (GrADS). This program integrates the functions of data access, data manipulation, and data visualization, providing users a single interface to their own data [Kinter and Doty, 1993]. The GrADS program has been designed to minimizing dependence on expensive computer hardware components such as system memory and graphics boards. GrADS runs well on most UNIX platforms (Sun Space, SGI, IBM, HP, DEC, etc.), as well as 80386(87)/80486/Pentium-based personal computers. We have introduced this program to an international community of geophysicists who have found it quite suitable for their interactive data analysis requirements. The following sections briefly describe the GrADS solutions to the problems outlined above.

2. USER INTERFACE

The GrADS program implements a command line interface in which the operations of data access, data manipulation, and data display can be performed by entering commands at a program prompt. To control data access, the user enters commands that describe the spatial and temporal

subsets of interest. To manipulate the data, the user enters expressions that specify the operations to be performed on the data. To control the display of the data, the user has a rich set of commands available to specify display types and attributes.

The GrADS user typically interacts with the program by first entering commands to open one or more data sets, then entering commands to set the desired dimension and graphics environments (or accepting the defaults provided by GrADS). Once this is done, the user may issue the `display` command with an expression as the operand of the command. This single command results in data being accessed, operated on, and displayed. When the user views this display and, as a result, wants to see something else, the user can simply enter another `display` command. If this is done without an intervening `clear` command, the new display is overlaid on the first. Judicious settings of the display environment used with overlaid displays can produce striking graphics with a high degree of quantitative information content (see Figure 1).

The GrADS scripting language (GSL) provides a programmable interface to the GrADS package [Doty and Kinter, 1993]. GSL is implemented as an interpreted language with a full set of operators and flow control capabilities. Terminal and file input/output is supported. The interface between GSL and the rest of GrADS provides a rich querying capability for data values, data attributes, the dimension environment, and the graphics environment. Users can develop scripts (using a standard text editor) to issue sequences of GrADS commands under program control. While executing the GrADS program, users can run the scripts they have created to allow a quasi-batch mode of operation (including a UNIX command line option to allow access to the shell) or to realize more complex or sophisticated data analysis and display capabilities.

The GSL incorporates a capability for users to develop a simple GUI to their own data. The GrADS command language includes commands for drawing simple widgets on the screen, such as buttons or sliders. Once the screen has been configured with combinations of widgets and graphics, a GSL script can query the location of the next mouse point-and-click. This location is returned to the script either as a widget number and associated values (if the click was on a widget) or

with coordinate values (if the click was on a graphic). When the user clicks on a widget, some appropriate action can be taken. When the user clicks inside a graphic, such as a contour plot, the coordinates can be transformed to world or grid coordinates and to some action taken, such as zooming or querying data. These simple capabilities provide for a wide variety of data GUI functions that users can design and implement in a very short time.

Several users have employed these basic building blocks to develop highly sophisticated graphical interfaces to their data. Examples of scripts that GrADS users have written include:

- A script that provides a point-and-click interface for viewing real-time surface airway meteorological data. The script allows the meteorologist to select the plot to be displayed via a menu of buttons. Each plot is a fairly complex overlay of several graphics displaying two or more variables in different ways, so that their associations are visually obvious and quantitatively determined. A click on the displayed graphic results in a zoom centered on that location or, in the case of a nested zoom, results in the display of a time series of the selected variable at the nearest station.
- A script that provides a point-and-click interface for viewing sea surface temperature (SST) over the Pacific Ocean. The script begins with a shaded contour display of a longitude(abscissa)-time(ordinate) section of SST with the selected latitude fixed. After pointing and clicking twice on the shaded contour plot, an animation is displayed of latitude-longitude contour plots, animated through the time period bracketed by the two clicks.
- A script that provides a point-and-click interface to the NOAA/NMC Nested Grid Model output. The script displays a menu of buttons that allows variables to be selected and plots for desired regions to be displayed. The menu also provides a time bar for selecting a specific time or a range of times for animation. The user may also display multiple variables by overlaying different graphics.

These scripts, and many others, have been developed by users to give them an interface to their data, rather than an interface to the GrADS program or to the graphics. We plan to support further development for such data GUIs.

The GrADS help facility is provided through a set of scripts that allow the user to review various commands and options via a point-and-click interface. Examples that demonstrate the use of various commands and options can be displayed graphically on the screen. In addition, a plain text manual describing the full set of GrADS commands and options is provided.

3. THE DIMENSION ENVIRONMENT

The concept of the *dimension environment* is fundamental to the GrADS program. The user describes a space and time subset of interest. This is specified as a combination of fixed and varying dimensions. The number of varying dimensions determines whether the subset is one-, two-, or three-dimensional. To specify a fixed dimension, the user would enter a single value for the dimension. To set a varying dimension, two values would be entered, indicating the desired range for that dimension.

The dimension values may be specified in either world coordinates or in coordinates specific to a particular data set (i.e., file coordinates). When file coordinates are used, they are converted to world coordinates using the transformation associated with a particular file, referred to as the *default file*. Even if we specify the dimensions in terms of file coordinates, GrADS maintains the dimension environment in terms of the world coordinates. This world view of the data is then translated back to file coordinates when data are actually accessed.

The dimension environment is a concept which is independent of the data. When we set the dimension environment, we are, in effect, describing a viewport into the four-dimensional data world. We do not need to worry about where those data are located in a particular file. For example, if we set the time at January 1, 1993, and open several files, we may find that January 1, 1993, is the first record in the first file, the 366th record in the second file, and the 32nd record in the third file. When data are accessed from any particular file, the January 1, 1993, record will be accessed, even though the user

has not explicitly referred to any particular physical location of the data in any of the files.

The spatial view of the data is also handled in a file-and-data independent fashion. For example, suppose that two files have been opened. The first is global in extent and contains a grid that is 1° by 1° in resolution. The second file has a 0.2° grid and covers only North America. To set the spatial viewport to the quadrant of the globe that includes North America, the following commands are entered:

```
set lon -180 0
set lat 0 90
```

When data are accessed from either file, we will get the data that cover that quadrant. In the case of the first file, those data are on a $1^\circ \times 1^\circ$ grid, and that grid is fully populated. In the case of the second file, the data are on a finer grid, and the areas in the quadrant outside North America (such as the central Atlantic Ocean) are undefined, because those areas do not exist in the second data file. The undefined areas of the grid are automatically filled with missing data values. If graphics from both grids are displayed, they are correctly overlaid with missing values treated appropriately.

It is important to repeat that the user need not be concerned about the physical storage particulars of the data. Data can be easily and quickly accessed and displayed from the perspective of the dimension environment, which provides a consistent geophysical view of the data that is integrated throughout the software.

4. DATA ACCESS AND FORMAT

The GrADS program is structured to access data from a file based on the current geophysical viewport, described by the dimension environment. Transparently to the user, GrADS translates world coordinates into file coordinates, computes the physical location of the data in the file, and accesses the desired subset of the data. Because the actual file structure and data format are decoupled from the user view of the data, any file structure or data format can be handled.

The particulars of a data file are described to GrADS via a *data description file*. This is a flat text

file that uses blank-delimited keywords to describe a particular data file. In general, the description file describes two aspects of the data. The first aspect is the orientation of the data in space and time. This provides for translating the user's view of the data into the physical storage organization of the data in the file. The second aspect of the file that must be described is the specific format of the data. This includes such things as byte ordering, record format, file headers, and so on. The data description file also contains the missing data value for the file. When encountered, this missing data value is assumed to represent data that is unavailable. This value is used throughout the entire GrADS program, so that missing data are handled properly during data operations and display.

The rationale for the data description file is that a given format must be describable to be useful. In this context, a describable file format is one whose characteristics can be written down in a simple finite form, such that an individual element of the file can be mapped into its spatial, temporal, and scientific domain (such as variable type and units). The GrADS program has been designed so that, given a data set with a describable file format, a GrADS description file paradigm can be implemented to describe the data format in the GrADS framework. The implementation of support for a describable format within GrADS is straightforward.

Currently, GrADS supports two gridded file formats, with a variety of minor variations allowed within the supported formats. The first format supported is the native GrADS format. This format is a simple IEEE floating point binary format, with the data elements ordered in a specified way. This format is designed so that it can be created, and later read and/or modified, by a simple FORTRAN or C program. The files are portable among all UNIX systems (including Cray UNICOS) and PCs. Minor variations to the native GrADS format allow the byte ordering of the IEEE words to be specified, allow the ordering in the Y and Z dimensions to be specified, and allow for file and record headers. The data may be written using direct access (stream) I/O, convenient for C programs, or using sequential I/O, convenient for FORTRAN programs.

GrADS also supports the GRIB records format established as a standard by the World Meteorological Organization. The GRIB records format packs

grids into an arbitrary number of bits per grid element and contains a specified record header that describes the data. In the case of GRIB, the data format is well described, but the translation between the ordering of data in a file and the orientation of the data in space and time cannot be predetermined. Thus, a GrADS utility is executed that creates an index file before GrADS is used to analyze and display the data in that file. This file is used by GrADS to quickly determine the location of a particular item of data in the data file, once the user request has been made.

We feel that the concept of indexing the data, to allow efficient and *user transparent* access to randomly organized data, is very important and powerful. In the case of GRIB, it allows the GRIB records to be ordered in any way in the file, and it allows partitioned records to be handled completely transparently to the user. Partitioned GRIB records occur when a single global grid is partitioned into multiple records so that a particular record does not exceed some arbitrary length.

In the case of station data, GrADS currently implements its own station data format. The GrADS station data format stores the data in an IEEE floating point binary format. The file is structured into reports, where each report contains a header (which describes the station identifier and the latitude, longitude, level, and time of the data) and one or more collections of data. Multiple levels of data may be contained in one report. In a way similar to the handling of GRIB data, a utility is run to index the station data.

GrADS can treat multiple data files as a single logical GrADS file, when the data files are split in the time dimension. For example, given monthly samples over a period of years, each month could be stored in a separate file, or each year could be stored in a separate file (12 months per file). GrADS determines which file contains what time periods by parsing the file name. The user specifies the file name using a naming convention expected by GrADS and specified as a file template in the data description file. The template specifies which temporal elements (year, month, day, hour, and forecast hour) are used and in what order they are used in the names of files to be treated as a logical unit. When a data access is requested, the date/time of that access is used to fill in the full file name.

The data are then read from that particular file. No more than one data file is kept open at a time, which avoids exceeding the open file limit that exists on most operating systems. If a particular file is missing, GrADS views it as missing data, and when displayed, it is handled as though the file exists but is filled with missing data values.

The file name template is particularly useful when handling real-time data. New files can be created and old files archived. Only the description file needs to be changed (to reflect the new time extent of the GrADS file), and utilities are provided to allow the time extent of a data description file to be updated “on the fly.” If the real-time data system should fail for a period of time, then the data files for that period would be missing. No special handling is required, because GrADS will automatically assume that the data is missing and handle it appropriately.

5. OPERATIONS ON DATA

Operations may be performed on data in GrADS by entering an expression. Expressions consist of combinations of operators, variables, and functions. A given expression operates over the range of the current dimension environment.

Operators within a GrADS expression are +, -, *, and /, and the operations are addition, subtraction, multiplication, and division, respectively. Operations take place over the entire range of the operands, which are variables, functions, or constants. An example of a GrADS expression is:

```
ts*9/5+32
```

In this example, a variable called *ts* is being converted from Centigrade to Fahrenheit. It is beyond the scope of this paper to document the entire syntax of the expression parser, but simple examples will be shown to illuminate specific points.

Variables represent “raw” data, which are contained either in a data file or in memory (a `define` command allows users to move data subsets to memory for more efficient manipulation). Variable names are specified by the user, either within the data description file (where GrADS variable names are assigned to data within a data file) or when the data are moved to memory.

The range of the data for a given variable is assumed, by default, to be the range of the current dimension environment. Within the expression, the user may *override* the current dimension environment. This allows, for example, a time difference of a variable to be easily calculated:

```
ts-ts(t-1)
```

In this example, the variable name is *ts*. A difference is being calculated between that variable at the fixed time specified by the dimension environment (the default) and the same variable at an offset from the time specified in the dimension environment (in this case, one time step earlier).

Multiple data files may be opened and variables accessed from any of those open files. An example of an expression that operates on data from different files is:

```
(ts.2-ts.1)*9/5+32
```

In this case, the variable *ts* is being differenced between two separate files (“2” refers to the second open file while “.1” identifies a variable from the first open file) before being converted to Fahrenheit.

Functions are also supported as part of the expression parser. Built in functions are provided to perform a variety of basic mathematical and statistical functions and to calculate simple meteorological variables, such as vorticity. Some functions can operate over a large range of data, such as the averaging function. Arguments to this function specify the dimension range over which the averaging is to be done. A simple nested expression could, for example, perform a temporal and zonal average over several data sets. Such an expression could process several gigabytes of data.

Another interesting function is *maskout*. This function takes two arguments. The operation of the function is such that the first argument is masked by the second and becomes the result. Whenever the second argument is negative, the corresponding data element in the result is set to missing. A typical application is to mask based on a land/sea grid. A *maskout* function, nested in an averaging function to calculate a global mean, could yield, for example, a time series of a variable averaged globally over only ocean points.

GrADS also supports functions written by users. Such functions operate externally to GrADS. These external functions are described via a *User Defined Function Table* (UDFT). This table is scanned when GrADS is first started, and functions described in the table are made available to the expression parser. When an expression is entered that uses such a function, GrADS goes through the following steps:

- The arguments to the function are evaluated. If the arguments are expressions, the expressions are evaluated and the results obtained.
- The value of each argument is written in a file in a documented format.
- The user defined function is invoked as a separate process. This program may do any processing desired, including invoking GrADS recursively, reading GrADS data files, and so forth.
- When processing is complete, the function writes its result into an output file, which is also in a documented format. The function then terminates execution.
- GrADS reads the result of the function (from the output file) and continues the evaluation of the expression.

Several functions are provided with the GrADS distribution as examples. One such function interpolates from any arbitrary grid spacing to a different grid spacing. This allows operations to be performed on grids that have different spatial resolution.

6. GRAPHICS OUTPUT

The result of a GrADS expression can be displayed using a variety of standard graphics output techniques. One-dimensional displays include line graphs and bar charts. Two-dimensional displays include contour plots (either level curves or shaded regions), wind vector or wind barb plots, streamlines, grid box values, and station model plots for station data. All vector displays (vectors, barbs, and streamlines) may be colored using any scalar field (which may itself be the result of a GrADS expression).

The user has a wide variety of options to control the way graphics are displayed. These options default to geophysically desirable values. Of particular note is the ease of control over the colors of contours, shaded contours, wind barbs,

wind vectors, and streamlines. The different graphics types may be easily overlaid to create complex graphics showing relationships between several different variables (Figure 1).

It is important to note that the dimension environment determines the range of the display. We will explain this concept with an example. Suppose the dimension environment is defined as follows:

```
set lon -180 0
set lat 40
set lev 500
set time 06dec1982 11dec1982
```

The Y and Z dimensions are fixed, while the varying dimensions are X and T. Thus, the axes of the output plot are X (longitude) and T (time) over the ranges specified (Figure 2). Annotation of the axes is automatic and appropriate to the dimensions selected for display. If the result of the expression reduces the number of varying dimensions (such as applying the averaging function), the display would be a one-dimensional plot, such as a line graph, where the varying dimension is X or T (Figure 3). Note that it is valid for an operation on data to reduce the number of varying dimensions, but it is invalid for an operation to increase the number of varying dimensions or to change which dimensions are varying.

Any graphics displayed on the screen can also be printed. The graphics are output in vector form to an intermediate metafile format, and they can then be translated to monochrome or color postscript using GrADS utilities that are provided with the GrADS program. Because the postscript files are vector graphics based, they will be rendered at the printer using the full resolution of the particular printer.

REFERENCES

- Doswell, J., *Forecaster Workstation Designs: Concepts and Issues*, *Wea. and Forecasting*, 7, 1-16, 1992.
- Doty, B.E. and J.L. Kinter III, The Grid Analysis and Display System (GrADS): An Update. *Proceedings of Ninth International Conf. on Interactive Info. and Processing Sys. for Meteor., Oceanogr., and Hydrol.* (Amer. Meteor. Soc.), 165-167, 1993.
- Kinter III, J.L. and B.E. Doty, The Grid Analysis and Display System (GrADS): A Practical Tool for Earth Science Visualization, *Info. Sys. Newsletter*, 27, 1-3, 1993.




Figure 1. Map displaying several variables at once. The level curves depict temperature with red tones for high values and blue tones for low values. The shaded contours represent relative humidity, with brighter green for higher values. Winds are depicted using standard meteorological wind barbs.




Figure 2. Longitude (abscissa) time (ordinate) diagram of 500-hPa geopotential height.




Figure 3. Line graph of 500-hPa geopotential height averaged by GrADS over a range of times at a fixed latitude displayed as a function of longitude. The open circles represent gridpoint locations, and the bars represent standard deviation (computed by GrADS) about the mean of the times being averaged.

Chapter IV

RELATED TOPICS

Making Sense From Space-Time Data in Laboratory Experiments on Space Plasma Processes

WALTER GEKELMAN, JAMES BAMBER,
DAVID LENEMAN, STEVE VINCENA,
JAMES MAGGS, and STEVE ROSENBERG

Department of Physics, University of California,
Los Angeles, CA

A number of visualization techniques are discussed in a laboratory experiment designed to study phenomena that occur in space. Visualization tools are used to design the apparatus, collect data, and make one-, two-, and three-dimensional plots of the results. These tools are an indispensable part of the experiment because the data sets are hundreds of megabytes in size and rapid turnaround is required.

1. INTRODUCTION

As a result of advances in general technology, it is now possible to perform laboratory experiments relevant to space plasma physics that were not possible 15 years ago. These advances include developments in computers, digitizers, and other hardware, as well as large improvements in the technology of plasma devices and plasma sources. The increased sophistication of laboratory techniques allows laboratory studies of the fundamental physics of space plasma processes. The relevance of laboratory studies to space plasmas is ensured by keeping constant ratios of dimensionless parameters. A successful laboratory study of a space plasma process involves selecting a particular phenomenon important to space plasmas that can be reproduced in a laboratory device in a scaled sense and performing measurements of its space and time evolution. Laboratory experiments can probe a process with unprecedented detail and uncover effects that may not easily be detectable in space.

Plasma sources now exist that can produce quiescent and highly reproducible plasmas. These sources have been developed to allow experiments in which highly detailed space-time measurements of electromagnetic fields and plasma parameters can be routinely collected. Understanding and control of the plasma boundaries have developed. Plasma properties, such as density gradients, striations, collisionality, magnetic field profiles, ion composition, and percentage of ionization, can be tailored for investigations of an individual process.

The development of laboratory plasmas alone would not allow the revolution in space laboratory experiments to occur. The accompanying strides in data acquisition [Gekelman, 1992] and workstations make possible the data collection and analysis

needed to process the experimental measurements. The advances in data acquisition technology include fast, and relatively cheap, workstations that have enormous amounts of ram and disk storage, fast (> 1 GHz) analog to digital converters, digital oscilloscopes, and programmable function generators that can be interfaced to computers. The result of applying sophisticated acquisition techniques to laboratory experiments is data sets in excess of 100 Mbytes. In a decade, laboratory data sets may exceed 1 terabyte. It is clear that without the proper analysis and visualization tools, understanding the experimental measurements would not be possible.

This paper points out how and why experimental data sets have grown over the last decade. We then illustrate the many ways in which a modern laboratory has become dependent on visualization in its everyday operation. Finally, we attempt to predict what the capabilities of a plasma physics experiment and its associated theoretical effort will be 10 years from now, as well as list the demands it will place on future visualization systems. As this book attests, there are many groups involved in plasma physics that use visualization techniques. Because this is not a review article, we cite examples from work done at UCLA as illustrations. The other papers demonstrate how widely visualization techniques are used in the space plasma physics community.

2. THE LAPD EXPERIMENTAL CONFIGURATION

Twenty years ago, nearly all plasma physics experimental data were in the form of x-y plots and tables. Digitizers, as we know them today, did not exist. Transient data were captured by photographing oscilloscope traces while sampled or continuous data were recorded with plotters. For example, if a waveform was recorded as a function of time at a fixed position, and the receiver was moved to other positions, the motion of a phase point could be laboriously measured using a ruler. When it is necessary to acquire two- or three-dimensional data, the data collection and analysis must be automated. We will illustrate this with data taken from the LAPD (Large Plasma Device) at UCLA. The LAPD [Gekelman et al., 1991] is a flexible, low-maintenance device designed to study a variety of

waves and nonlinear effects in fully magnetized plasmas. The plasma column is 50 cm in diameter and 10 meters in length. The magnetic field is controlled by seven independent supplies, which allow tailoring of the axial magnetic field profile. The maximum DC axial field is 2.5 kG which gives a highly magnetized plasma (a He plasma column is 500 ion gyroradii wide).

The plasma is produced by a DC discharge driven by an oxide-coated cathode. This type of plasma has proven quiescent ($\delta n/n \approx 5\%$), and oxide-coated cathode sources are stable for long periods (≈ 6 weeks), giving plasma discharges that are very reproducible from shot to shot. The plasma (H, He, Ar, Xe, Kr, Ne, or any mixture) can be moderately high in density ($n \leq 5 \times 10^{12}/\text{cm}^3$), and at the LAPD laboratory, we have achieved very highly ionized helium plasmas, with the percentage of ionization about 99 percent [Maggs et al., 1991]. The discharge is pulsed at several Hz to allow for efficient signal averaging and data processing. The plasma is diagnosed using a completely internal probe drive capable of moving, under computer control, an assemblage of magnetic and electric field sensors to any position within the plasma volume [Pfister et al., 1991].

A volume data set is acquired by measuring the three components of the electric and magnetic fields at up to 32K time steps at a given location. Langmuir probes are used to measure the plasma density, plasma potential, and electron temperature at up to 16 times during a shot. The Langmuir probe is rapidly swept, and the probe current and voltage are digitized. Housekeeping functions, such as the neutral gas pressure and cathode temperature, are also digitized and stored. The probe is moved to the next position in the volume, and more data are taken. The data run could store results at 10K positions, resulting in data sets of 300 Mbyte. If experimental conditions are changed, different data sets are generated, each of which must be analyzed rapidly so that interesting features can be isolated and future data runs effectively planned.

Of paramount importance in a well-diagnosed basic physics experiment is a state-of-the-art data acquisition system. The LAPD laboratory has assembled hardware and software that enable real-time storage and analysis of very large data sets [Gekelman and Xu, 1986]. The system is flexible

enough to perform real-time manipulations of incoming data and provide rapid turnaround for data reduction and display. The front-end of the system consists of a VAX station 3900, which controls both CAMAC and GPIB systems in which analogs to digital converters, programmable amplifiers, stepping motor controllers, and so on, are located. The acquisition system is, in turn, networked to a Kubota Titan Super Graphics Workstation, DEC station 5000's, and an SGI 240/VGX. These computers have hardware that can render three-dimensional images and rotate and zoom in real time. This capability is indispensable for visualizing the complex data sets generated from the types of experiments conducted at the LAPD laboratory. The workstations have up to 128 Mbyte of RAM and a total of 9 Gbytes of disk storage. They run the latest in visualization software, such as AVS, Data Visualizer, PV-Wave, and Advanced Visualizer. In addition, our group has written many data analysis programs, as well as all the data acquisition software.

The local laboratory network is linked with fiber optics to the rest of the UCLA campus, to access other research computers, and to a large variety of external networks. The most important link connects the LAPD with the UCLA Visualization Center, which was constructed, in part, by one of the authors (Gekelman). The key piece of equipment in the center is an Abekas A-60. This device will store 50 seconds (1,500 NTSC frames) of digital video; data can be loaded to across the network. The Abekas may be programmed to play the data back in segments onto a Sony Betacam-SP video recorder. The segments can be of any length and played back at any speed. The Betacam is used for a master recording. Video movies may then be transferred to VHS or Super VHS in any format (NTSC, PAL, SECAM). This computing power is indispensable for visualizing the complex data sets generated in the modern plasma laboratory. In the past year, we have come to depend on making videos quickly (overnight) to understand the data and alter ongoing experiments. Some of the real-time aspects of this analysis will be presented here. Our experiment and analysis tools have become so interdependent that we consider visualization an integral part of our laboratory.

3. GRAPHICS PACKAGES

The appearance of flexible and easy-to-use graphics packages is a recent phenomenon. Fifteen years ago, people wrote their own graphics packages, which were highly device dependent. Graphical output was often screen dumps from the terminal and line printer output. For want of video technology, 16-mm movies were made to illustrate three-dimensional effects. Because of the difficulty involved in their production, movies were made only for presentations at scientific meetings. For example, a movie made by one of the authors (Gekelman) showing the tearing of a current sheet [Gekelman et al., 1989] took about 2 months to complete. The process involved passing the data through a public domain three-dimensional contouring program (GRAPE), translating it from VAX 32-bit format to CRAY 64-bit, and inputting the contours to a surface builder and renderer written at the San Diego Supercomputer Center. Vector fields were treated with an algorithm written at UCLA, which drew three-dimensional vectors on the rendered data. The project consumed 120 hours of CRAY-YMP time and took 2 weeks to render. The final product was high-quality 16-mm film, but once a 16-mm film is made, there is no way to change it.

That example is in marked contrast to what can be done today. Data are collected and moved across the network to a graphics workstation where they are translated from VAX to UNIX binary in a matter of minutes. The data are then rescaled and transformed from the probe to the laboratory coordinate system. A data set can be ready for rendering in several hours. A 1,500-frame video is programmed (selecting views, titling, and colorizing) in an hour and can be rendered overnight on one of our graphics workstations. It takes about 2 hours to transfer the frames to the Visualization Lab Abekas (this is an ethernet network bandwidth limitation) and another 10 minutes to make the video. In the future, we envision the entire process to be limited only by the time it takes to choreograph the video!

Table 1 lists some of the graphics packages used by the authors, as well as the source of the package and its principal use. There are, of course, other packages.

Table 1. Graphics packages used by the authors.

Package	Author	Symbols	Use
AVS	AVS INC.	\$\$,N,I,T,A,M,nw,U,RT	3-D, 2-D, data analysis, network based
Data Visualizer	Wavefront Technologies	\$\$,N,I,T,A,m,nw,U,r,RT*	3-D, 2-D data analysis, very easy to use
Advanced Visualizer	Wavefront Technologies	\$\$\$,I,nw,A,r,RT,T	3-D, model building, animation standard for film industry
Photoshop	Adobe	\$.I,u,T	image processing, post production
Vellum	Ashar	\$.I,m,T	2-D, 3-D drafting
PV-Wave	Visual Numerics	\$\$,N,I,M,T,A, nw,U,r	1-D, 2-D fast data analysis, contouring
TV80lib	NERST	Pd	1-D, 2-D wireframe, contours
movie BYU	Brigham Young University	Pd,T	1-D, 2-D surface, contours
NeoVisuals	Neovisual Inc.	\$\$,RT	first decent ray tracer
Snapshot	Silicon Graphics	N,I,T	paint program with image processing
SDCS Image Tools	SDSC	PD,N,I	image translator
Pandemonium	XAOS Tools	\$\$,T,RT	ultimate titling
Doré	Kubota Pacific Inc.	\$.I,T,A,r	real-time 3-D animation

Key to Symbols

\$ = moderate in cost less than \$1K

\$\$ = expensive less than \$4K

\$\$\$ = very expensive more than \$4K (sometimes up to 80K)

Pd = public domain, usually price of media

N = new, less than 5 years old

nw = can be run over a network

RT = has a ray tracer

r = renderer (Gourard or Phong) *Some programs such as AVS and the Data Visualizer take advantage of the host computer's specialized hardware.*

I = interactive

m = has some mathematical functions

M = has extensive mathematical functions

T = has some titling capabilities

A = can make animation

u = user may add own modules

RT* = will use ray tracer of the Advanced Visualizer (if you have it)

SDSC = San Diego Supercomputer Center

4. EXPERIMENT SUPPORT

The remaining part of this article illustrates how visualization tools are used in almost every step of our laboratory's operation. CAD tools are routinely used to prepare drawings that can be submitted to machine shops. We have used Vellum; it is very easy to use, and we make high-quality drawings of everything we build. Previously, we were content with sketches.

Apart from designing instruments, visualization tools are sometimes essential in understanding their measurement capabilities. One example of this is in describing the transmission function of a directional velocity analyzer. These devices are constructed by interfacing a collimator (array of narrow channels) and a conventional velocity analyzer [Stenzel et al., 1982; Stenzel et al., 1983;

Leneman et al., 1991]. The behavior of this instrument is straightforward in an unmagnetized plasma, but difficult to describe in a magnetized plasma because of the helical nature of the particle orbits. To understand the operation of the directional velocity analyzer in a magnetized plasma, a computer program was written to calculate the instrument sensitivity as a function of perpendicular and parallel particle velocity for a given species (ion or electron) and background magnetic field strength. The numbers generated are then rendered as a surface whose height above a point in the Cartesian plane represents the sensitivity $S(v_{||}, v_{\perp})$ of the analyzer. The surface shown in Figure 1 was colored to help easily identify changes in the sensitivity. The axes were added afterwards using Photoshop on a Macintosh Quadra.




Figure 1. Instrumental sensitivity, $S(v_{\perp}, v_{\parallel})$, of a directional velocity analyzer shown as a function of the components of the velocity of the detected particle perpendicular and parallel to the local magnetic field. The red areas show regions of heightened sensitivity, and the blue areas show sensitivity too low to be useful experimentally. Black represents the region in velocity space from which no particles can enter the velocity analyzer.

5. TWO-DIMENSIONAL DATA

Two-dimensional data may be rendered as contour maps or color-coded surfaces (or a combination of the two). As an example, we show data from an experiment on whistler waves. Figure 2a shows wave phase fronts in a plane that contains the background magnetic field. The plasma density is $n = 2.0 \times 10^{11}/\text{cm}^3$, the whistler frequency is $f_{\text{wave}} = 80 \text{ MHz}$, and the ratio of the wave frequency to the electron cyclotron frequency is $f_{\text{wave}}/f_{\text{ce}} = .07$. The red areas are wave magnetic field maxima, and the blue are minima.

Generally, those involved in an experiment know where all the relevant probes, antennas, and boundaries are located. However, when communicating results to others not familiar with the geometry, the ability to render relevant components of the apparatus on actual data is extremely helpful. This is illustrated in Figure 2a, which depicts the exciter at the left, and Figure 2b, which shows the whistler data as it is located in the device. The orange square at the rear is the oxide-coated cathode, which is the plasma source. The circular lines

show the boundary of the vacuum vessel. The exciter and data plane are in the midground. The non-data objects may be built by calculating all points on their surface and giving them a single scalar value; AVS then renders them as an isoscalar. Very complicated objects can be built with the Advanced Visualizer and then exported to AVS or the Data Visualizer. We have found that, when the data are complicated, placing objects in three-dimensional renderings is useful. Obviously, these techniques will soon be used for spacecraft and satellite measurements, where complicated boundaries and measurement geometries exist.

Integrating three-dimensional data with digitized images is also useful for presenting the perspective of scale sizes. A cutaway view of the device with magnetic field data, B_y , placed within it is shown in Figure 3. The wave magnetic field shown is that of a shear Alfvén wave. Data were acquired on 10 parallel planes. The wave maxima are colored red, and the minima are blue. Green areas have zero wave magnetic field. One can see that B_y is largest in a field-aligned filament. The wave data were rendered with the Data Visualizer,




Figure 2a. Data of whistler wave magnetic field intensity is shown as a function of position. The waves were launched in a uniform plasma from a trombone-shaped antenna, which has been rendered in the drawing on the left. Superimposing the exciter helps orient the viewer. The broad wave pattern reflects the shape of the exciter.




Figure 2b. A plane of wave data is rendered, to scale, with important components of the experimental device. Shown is the exciter (in white), the cathode (the orange square at the rear with a darker orange, circular emitting area), and the walls of the device (white circles). A white measurement grid also is superimposed.




Figure 3. A cutaway view of the LAPD with measured Alfvén waves in it. The Data Visualizer was used to generate a series of 10 cutplanes, on which one component of the wave magnetic field (B_y) is rendered. Red is used to denote wave maxima and blue minima. The image was imported to Photoshop, along with an image of the device output from a slide scanner. A part of the machine slide was cut away and the data image superimposed. This technique allows those unfamiliar with the geometry of an experiment to become instantly acquainted with it.

and Photoshop was used to process the digitized image and superimpose the data.

These same techniques are indispensable to theoreticians. Figure 4 shows a two-dimensional surface representing the spatial evolution of the magnetic field-aligned velocity distribution of electrons interacting with the electric field excited near plasma resonance in a non-uniform plasma. The initial electron distribution consists of a low-temperature spatially non-uniform part and a hot uniform Maxwellian tail population. The surface is color coded and textured to help identify features. Red represents high particle density and blue low particle density. The high-energy tail population interacts with the resonant electric field to form a beam-like velocity distribution through a process of phase-independent acceleration [Maggs and Morales, 1990].

The surface was produced using AVS with data generated by a computer calculation of the velocity distribution using an array of test particle orbits. Representation of the data in this fashion reveals

details of the spatial evolution and helps illuminate the process of beam formation. Such a process may produce bursts of superthermal electrons in the Earth's ionosphere and is a candidate for experimental study in the LAPD.

6. VISUALIZATION TOOLS IN REAL-TIME DATA ACQUISITION

In laboratory experiments that generate volume data, real-time feedback between the data and experiment is a necessity. It is disappointing, as well as wasteful, to spend days acquiring an enormous data set that is uninteresting. Furthermore, unexpected effects (which plasmas seem quite willing to provide) sometimes result in having the effect of interest spill out of a preprogrammed measurement volume. A plasma wave experiment involves varying one or more parameters in a discrete manner and making measurements at each point in the parameter space. The spatial dimensionality of the measurements may range from a




Figure 4. The theoretically calculated velocity distribution, $f(v)$, of electrons interacting with an electric field at plasma resonance in a non-uniform plasma is shown as a function of velocity and space. Red indicates high density and blue low. The distribution starts out as a Maxwellian, but evolves into a distribution with several beam-like features after interacting with the wave. The beam-like features are indicated by the local maximums in the distribution on the surface nearest the viewer.

single point (zero-dimensional) to a full volume (three-dimensional). A three-dimensional grid on which data are acquired can have as many as 10,000 positions. The fixed parameters for a given run include the wave-launching antenna position, wave frequency, average plasma density, background magnetic field, and so on. Variables taken at each point can be the local magnetic field, electric field, plasma density, and electron temperature.

Measurements consist of digitized signals from probes inserted into the plasma. It is a simple matter to display these signals as a function of time as they come in. However, one is more interested in the variation of the data as a function of the fixed parameters and of position and time. Quantities such as plasma density or electric field must be extracted from the digitized signals. In general, significant processing is required to extract the quantities of interest. Basic processing includes offset removal, coordinate transformation, and other simple mathematical operations. More complex analysis includes curve-fitting, digital filtering,

correlations, interpolation, and so on. Statistical analysis also is sometimes desired. Then, depending on the dimensionality of the parameter space and the spatial dimensionality, the data will be displayed using the appropriate graphics package.

To avoid slowing the data acquisition computer, all of the post-processing runs independently on two separate workstations. One is devoted to analysis, the other to visualization. They share access to a large (~ 1 Gbyte) region of disk memory. In an experiment involving the interaction of a whistler wave with a density striation, the electric and magnetic fields, as well as the plasma density, are measured for a given incident whistler wave frequency. The measurements are made on a series of planes that are perpendicular to the background magnetic field. The data acquisition system begins by setting the frequency and moving the probe head to the first position in the data plane. Each data component is digitized and stored for a number of plasma shots (typically 10). After all shots have been taken, the data acquisition system updates a

status file, which contains information such as the name of the current data set and the number of data records it contains. Then it moves the probe to the next position, and acquisition continues. When the plane has been mapped, the next wave frequency is set, a new data file is opened, the probe is moved back to the first point in the plane, and the process is repeated.

Concurrently, the analysis process, which runs on a separate workstation, periodically checks the status file; if it finds that new data have arrived, the workstation proceeds with the analysis in the following manner. New data records are appended to a backup copy of the current data set. Then they are split into electric field, magnetic field, and plasma density components. The calculation then proceeds on two separate tracks. The current-voltage response, as measured by a Langmuir probe, contains density, potential, and electron temperature information. To extract these from the digitized Langmuir probe data, a curve-fitting procedure must be used. The electric and magnetic wave field data are analyzed quite differently. For each component of these data, the wave amplitude and phase are calculated from correlations with the input signal. When both types of calculation are complete, the median values of plasma density and the wave amplitudes and phases are stored. The median values are appended to a display file, and the calculation status file is updated to signal the visualization process that there are new data to be displayed.

The Application Visualization System (AVS) is used to display the data. To use AVS, one connects a number of simple program modules together to form a flow network. One module periodically checks the calculation status file for the existence of new data. If new data exist, other modules in the network are signaled to update the display. One component of the wave data is color coded and plotted as a two-dimensional surface. A contour of plasma density, which reflects the position of the density striation, is superimposed. The manager module also can be used to change which component is being displayed, and change data sets (i.e., different input frequencies), contour level, and so on. The time variation of the wave fields can be reproduced by advancing the wave phase in small steps. No matter how the display parameters are

changed, the picture is continually updated with newly analyzed data.

It takes on the order of 1 minute to accomplish all the probe motions and acquire and store all the shots of data at a single spatial location. Acquiring data over a full plane with decent resolution (30 by 30 positions, for example) requires approximately 15 hours. The analysis takes about 10 seconds per point, for a total of approximately 2.5 hours. Real-time display of data is critical for the experimenter to be able to make adjustments to experimental parameters and correction to the experimental setup in a reasonable amount of time. It is also extremely useful as a tool to enable a deeper understanding of the physics being measured and to guide the decision as to the experimental directions to be pursued.

7. THREE-DIMENSIONAL DATA

Data from plasma physics experiments usually consist of vector fields $[\mathbf{E}(\mathbf{r},t), \mathbf{B}(\mathbf{r},t)]$ and scalar fields $[n(\mathbf{r},t), T_e(\mathbf{r},t), V_p(\mathbf{r},t)]$. Three-dimensional scalar fields may be plotted as isosurfaces (i.e., surfaces on which the quantity has a fixed value). Vector fields may be plotted as arrows. Sometimes, a large array of three-dimensional vectors can be confusing. One can then combine both techniques to plot a vector field. This is shown in Figure 5, which is experimental data of the tearing of a current sheet [Gekelman and Pfister, 1988]. In this work, an initially uniform slab of electron current was forced to flow in a magnetoplasma. When the current sheet was much longer than its height, it tore into a three-dimensional array of magnetic islands and X type neutral points (locations at which the transverse magnetic field vanishes). To avoid confusion, the axial component (j_y) of the current was represented as a series of isosurfaces of varying transparency and color. These surfaces were parameterized by the axial current density. The transverse current was represented as a vector (j_x, j_z) plot. In this experiment, there was a background magnetic field in the y direction. One sees that structure has developed in the axial current. The locations at which j_z is zero occur at magnetic X points.

Finally, there is a class of three-dimensional data that a single image cannot capture. This is illustrated in a rendering (Figure 6) of two magnetic




Figure 5. Three-dimensional surfaces of axial current density J_y (orange - 1.2 A/cm², green 1.0 A/cm², and blue 0.5 A/cm²) are shown at $t = 4.75$ ms after the formation of a thin (width/length = 0.1) current sheet. The magenta arrows indicate current flow perpendicular to the background magnetic field. The yellow rectangle in the foreground shows the position of the slot through which the current is funneled at $t = 0$.

flux tubes that have become intertwined. The experiment in which this occurred involved the interruption of a plasma current [Gekelman et al., 1987]. The initial magnetic field lines caused by the plasma current are circles surrounding the current channel. When the current is forced to break up into eddies and return currents, the structure of the field lines and their associated flux tubes become complex. The two flux tubes shown here, arbitrarily colored red and blue, are linked in a complicated way. Unless the image is printed as a stereo pair (by plotting two images as seen by observers six degrees apart—the average angular separation of our eyes), it is not possible to determine whether the flux tubes are linked. Such a stereo pair was printed in the reference cited along with another flux tube pair, equally as complex but not linked.

Without the use of stereo viewing, there is no way to tell them apart. An increasing number of workstation vendors supply hardware that allow their systems to show color stereo images. This is accomplished by alternately placing each of the image pair on the screen and synchronizing electro-Polaroid glasses with it. The stereo viewing process is greatly aided by machines fast enough to render images in real time as they are rotated. In some cases, a non-stereo system that can rotate isosurfaces in real time may suffice.

8. DEVELOPMENTS IN THE NEAR FUTURE

Ten years ago, visualization tools were used mainly to present results to scientists at meetings and seminars. With the increase in diagnostics and




Figure 6. Two flux tubes, created when a plasma current flowing along a background magnetic field is disrupted by switching on a strong orthogonal magnetic field, are shown in blue and red. The magnetic flux tubes are initially circular and surround the undisrupted current. After the current is interrupted, the flux tubes become complicated and intertwine. A stereo view is necessary to determine whether the flux tubes become linked.

digitization techniques, they are now an integral part of the running experiment. Graphics are used to display data as it is taken and for rapid analysis of large data sets at intervals within an experiment. Images that used to be time and labor intensive to create and were only reserved for figures in publications are now used as graphical scratch pads and “quick looks.”

One of the limitations of probes in laboratory plasmas is their size. Because they are many Debye lengths (and sometimes electron Larmor radii) across, they perturb the plasma they are measuring. Laser diagnostics are non-perturbative but, because of expense, have been limited to only a few channels of data acquisition. This is changing. Fabrication techniques will enable the production of

microscopically small probes. The price of lasers is dropping, and inexpensive tunable solid state lasers are emerging. Along with these advances, there has been a steady drop in the cost of high-speed analog to digital converters. It is not hard to picture future laboratory experiments with hundreds of thousands of data channels generating up to 100 terabytes of data per run. Such experiments will be necessary to investigate chaotic and turbulent systems in a systematic way. These problems will be attacked by workstations of the future that will be armed gigabytes of ram and terabytes of disk or other inexpensive storage. It is obvious that visualization tools and scientific data base management techniques will be indispensable for every step of the process. Coupled to these hardware advances will

be innovative ways to present data. This could be done, for example, by allowing others to access animations of results over high-speed networks. Eventually, cellular technology will enable these networks to export films and journal articles to laptop computers. Today's means of presenting scientific information will soon become obsolete.

Acknowledgments. This work was sponsored by the Office of Naval Research under grant ONR N00014-91-J-1172 and by the National Science Foundation under grant NSF-ATM-9214000.

REFERENCES

- Gekelman, W. and L. Xu, Flexible, Real-Time, Data Collection and Analysis System and its Utilization in a Basic Plasma Physics Experiment, *Rev. Sci. Instrum.*, 57, 1851-1855, 1986.
- Gekelman, W., J.M. Urrutia, and R.L. Stenzel, Measurement of Magnetic Helicity During the Disruption of a Neutral Sheet Current, *Magnetotail Physics*, edited by Anthony T. Lui, Johns Hopkins Press, Baltimore, 261-274, 1987.
- Gekelman, W. and H. Pfister, Experimental Observations of the Tearing on an Electron Current Sheet, *Phys. Fluids*, 31, 2017-2025, 1988.
- Gekelman, W., H. Pfister, and R.L. Stenzel, Experiments on Magnetic Field Line Reconnection in Two and Three Dimensions, *Trans. AGU*, 15, 426, 1989.
- Gekelman, W., H. Pfister, Z. Lucky, J. Bamber, D. Leneman, and J. Maggs, Design, Construction and Properties of the Large Plasma Research Device—The LAPD at UCLA, *Rev. Sci. Instrum.*, 62, 2875-2883, 1991.
- Gekelman, W., Data-Acquisition Systems, *Encyclopedia of Applied Physics*, 4, 464-483, VCH Publishers, 1992.
- Leneman, D., W. Gekelman, J. Maggs, H. Pfister, and J. Bamber, Improved Techniques in Velocity Analysis, *Bull. APS*, 36, 2327, 1991.
- Maggs, J.E. and G.J. Morales, Nonlinear Dynamics of Electrons Accelerated by Resonant Fields in Nonuniform Plasmas, *Phys. Fluids B*, 2, 3134-3148, 1990.
- Maggs, J.E., R.J. Taylor, and W. Gekelman, Neutral Burnout in a Ten Meter Liner Device, *Bull. Amer. Phys. Soc.*, 36, 2415, 1991.
- Pfister, H., W. Gekelman, J. Bamber, D. Leneman, and Z. Lucky, A Fully Three-Dimensional-Movable, 10-m-long, Remotely Controllable Probe Drive for a Plasma Discharge Device, *Rev. Sci. Instrum*, 62, 2884-2889, 1991.
- Stenzel, R.L., R. Williams, R. Agüero, K. Kitazaki, A. Ling, T. McDonald, and J. Spitzer, *Rev. Sci. Instrum.*, 53, 1027, 1982.
- Stenzel, R.L., W. Gekelman, N. Wild, J.M. Urrutia, and D. Whelan, Directional Velocity Analyzer for Measuring Electron Distribution Functions in Plasmas, *Rev. Sci. Instrum.*, 54, 1302, 1983.

Envision: An Interactive System for the Management and Visualization of Large Geophysical Data Sets

K.R. SEARIGHT, D.P. WOJTOWICZ, J.E. WALSH

Department of Atmospheric Sciences, University of Illinois, Urbana, IL

S. PATHI and K.P. BOWMAN

Department of Meteorology, Texas A & M University, College Station, TX

R.B. WILHELMSON

Department of Atmospheric Sciences and National Center for Supercomputing Applications, University of Illinois, Urbana, IL

Envision is a software project at the University of Illinois and Texas A&M, funded by NASA's Applied Information Systems Research Project. It provides researchers in the geophysical sciences convenient ways to manage, browse, and visualize large observed or model data sets. Envision integrates data management, analysis, and visualization of geophysical data in an interactive environment. It employs commonly used standards in data formats, operating systems, networking, and graphics. It also attempts, wherever possible, to integrate with existing scientific visualization and analysis software. Envision has an easy-to-use graphical interface, distributed process components, and an extensible design. It is a public domain package, freely available to the scientific community.

1. INTRODUCTION

Envision is an interactive software package for the analysis and display of measured and modeled data sets [Searight et al., 1993a; Searight et al., 1993b]. The key features of Envision are support for data in NetCDF and HDF files, an easy-to-use X/Motif user interface, distributed process capabilities in a client-server arrangement, and portability to many UNIX workstations. The Envision package also provides the user new ways to view and change metadata in a set of data files. It permits a scientist to manage conveniently and efficiently large data sets consisting of many data files. It also provides links to popular visualization tools so that data can be quickly browsed.

1.1 Motivation

Earth science researchers work with increasingly larger data projects. They work with both large observational data sets and large model outputs. The Earth Observing System Data and Information System (EOSDIS) illustrates the magnitude of the data explosion in the earth sciences. EOSDIS will produce a large variety of data sets, ultimately containing more than 10 petabytes (10^{15}) of data (M. Folk, NCSA, personal communication, 1993). To carry out successful research on a project such as this, scientists need to be able to browse, visualize, and analyze large and varied earth science data sets.

There are a number of problems to address in designing software to work with large data sets. When dealing with large data sets, the issue of data management becomes critical. Another important consideration is the handling and use of metadata, or data about the data. Earth scientists have to deal

with data sets composed of many files and need interactive, random access to any part of them as transparently as possible. Data subsets should be easily and quickly displayed with a range of visualization options for analysis purposes. Researchers must also be able to integrate their own visualization and analysis programs with a data management system. These concerns form the basis for the design and implementation of the Envision package.

1.2 Project Goals

Envision was created to address the lack of effective tools to manage and browse large scientific data sets. Clearly, there are other software packages in existence that offer partial solutions to these problems. So, rather than starting from scratch, Envision was designed to take advantage not only of existing standards in file formats, hardware, and software, but also to interface with existing packages and software libraries, where possible. Envision's goal is to integrate the areas of data management, analysis, and visualization together in an interactive environment. The project's design is modular and distributed to allow for flexibility and extensibility, but it also tries to hide from the user unnecessary detail about the storage of data values in files. To achieve wider use, Envision is a public domain package that will run on most UNIX workstations.

2. SYSTEM ORGANIZATION

2.1 Requirements

The Envision package is designed to work with commonly used standards in data formats, operating systems, networking, and graphics. It runs on many UNIX workstations using X/Motif, including IBM RS6000, Sun, HP, and SGI. No special hardware is needed. Additional ports of Envision to other UNIX platforms are in progress.

Data sets intended to be used in Envision are stored as rectangular grids of values in n -dimensions. Envision currently works with the NetCDF format [Rew and Davis, 1990], developed by Unidata at the University Corporation for Atmospheric Research (UCAR) and HDF [Brown et al., 1993], a format developed by National Center for

Supercomputing Applications (NCSA) at University of Illinois. HDF is the format recently adopted by the ESDIS project as the baseline standard for EOSDIS data product generation, archive, ingest, and distribution.

2.2 Major Components

Envision's key features are a metadata browser and editor, a data management system, and a set of links to visualization and analysis tools. Figure 1 shows a schematic of the system, which is modular in design. Envision's principal components are the Envision Data Manager (EDM) and Envision User Interface (EUI), which run as separate processes in communication with each other. At the center of Envision is EDM, which acts as a server, to which any number of client processes, including EUI, may connect. Visualization and analysis tools are connected as clients, and EDM exchanges messages and data with them. Visualizations are performed using the public domain XImage and Collage programs from NCSA and the commercial IDL package from Research Systems, Inc., of Boulder, CO. The modular design of Envision provides extensibility, because other tools may be added to the Envision environment. EUI is an intuitive point-and-click graphical environment and has flexible customization features to view data sets.

3. DATA MANAGEMENT

Data management is one of the key purposes of Envision. EDM provides mechanisms to manage data variables, dimensions, and attributes in a group of related data files.

3.1 Managing Multiple-File Data Sets

Data used in a large project are invariably contained in a set of data files. In practice, this organization of data can be unwieldy to manage. One way a data set might be stored is on a CD-ROM, with data files stored in a hierarchy of directories representing years, months, and days. The data in these files are conceptually a single entity, but have been divided for convenience of storage or because of its collection over a period of time. In a scenario like this, problems would arise

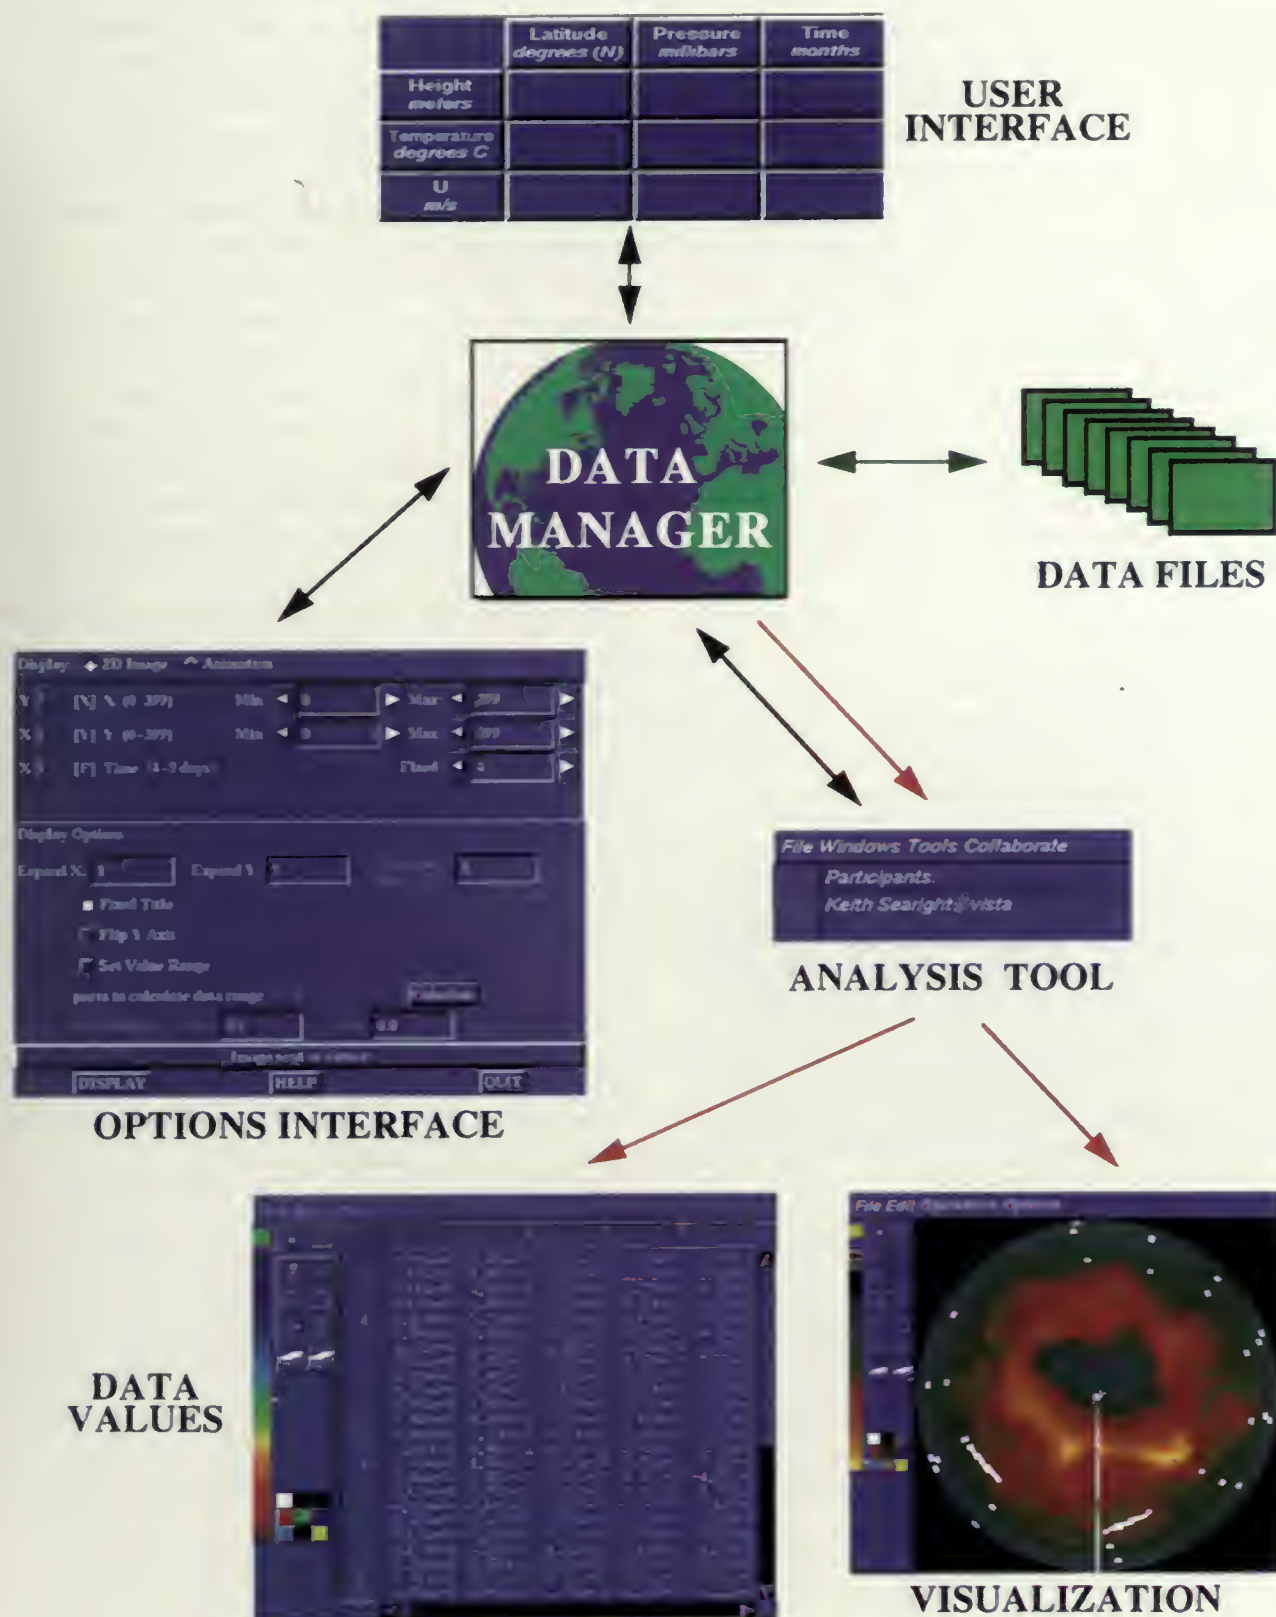


Figure 1. Envision schematic. Black arrows indicate passing of messages and metadata between separate processes. The green arrow indicates reading and writing of data values and metadata to and from files. Red arrows indicate direct transmission of data values by interprocess communication.

if a user wished to view a subset of the whole data set when the values wanted were stored in a number of these files. For example, a user might wish to see a display where time is an axis.

Management of variable and dimension objects that span multiple files is done in Envision by matching up the objects that correspond in each of the files. Envision organizes large data sets into projects and the metadata, or data about the data, is saved in project files. The operation of matching up and grouping the same variables and dimensions in different files is referred to as merging. This operation creates new entities, called logical objects. Logical objects are simply a way to group together regular dimensions and variables that span multiple files, so they can be used as if they were

all in a single file. Envision also handles situations where dimensions may be implied by the division of data into files. A dimension, quite often time, may not be explicitly defined, but can be assumed when each file in a data set represents a single step. In Envision, this type of dimension may be created and is called a virtual dimension. Once a virtual dimension is defined, it can be used exactly like any other dimension. Typically, these are created for each data file and then merged together to make a logical dimension.

A simple example will serve to illustrate how the merging of dimensions and variables into logical objects works. Figure 2a shows a set of three data files, each containing one variable, TEMPERATURE, and the two dimensions it uses,

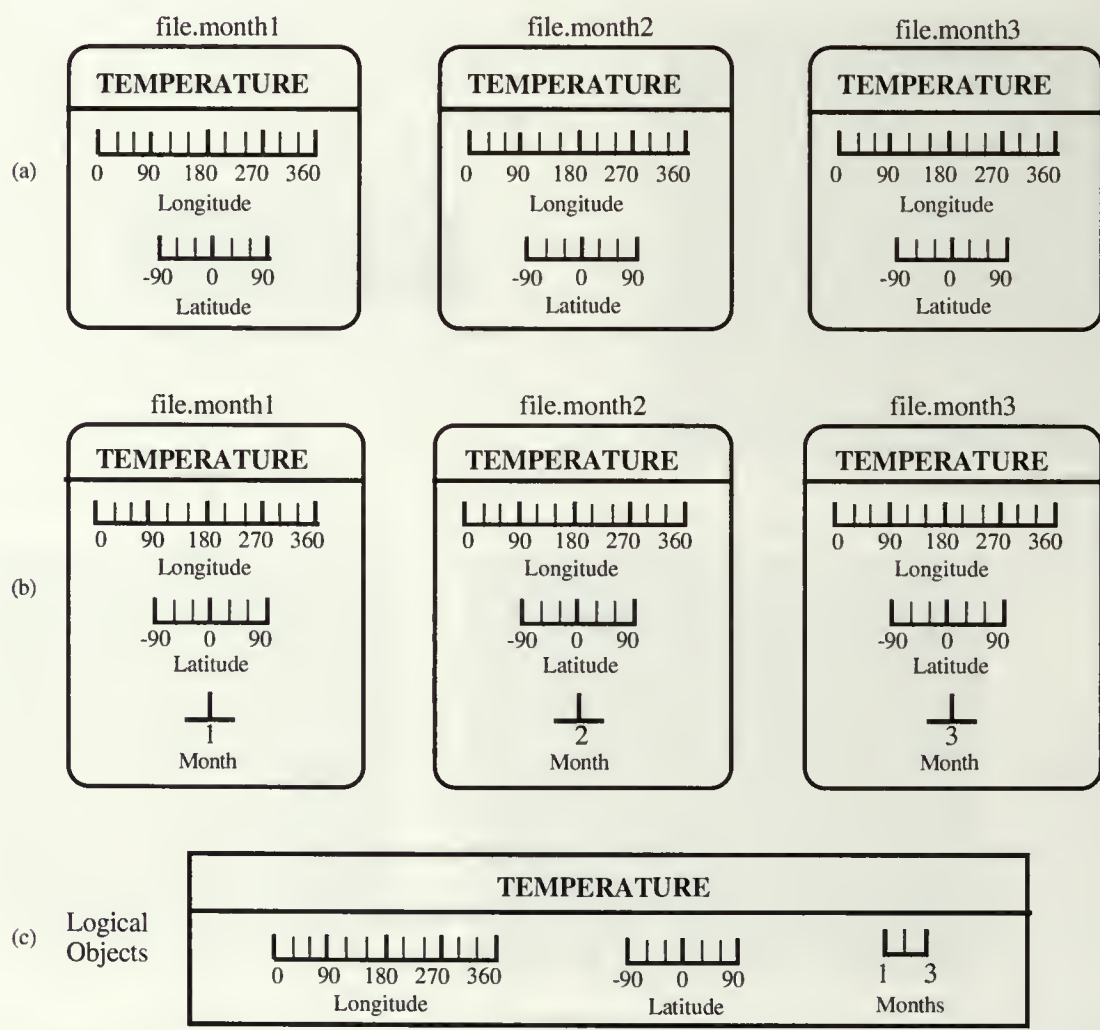


Figure 2. Creation of logical dimensions and variable using a virtual dimension—(a) three related data files, each representing a 1-month time step; (b) the same data files after creating the virtual dimension, Month, and adding it to the variable; (c) logical dimensions and variable after merging.

Latitude and Longitude. In each file, the Latitude and Longitude dimensions have coordinate variables with the same set of values as shown. There is no time dimension specified, but each file represents a single month time step. The variable, TEMPERATURE, contains $7 \times 13 = 91$ values in each of the data files. Figure 2b shows the same three files after a virtual dimension, Month, has been created for each and then added to the variable, TEMPERATURE. These changes are made only in the Envision project, and the data files are not actually modified. Note that TEMPERATURE has $7 \times 13 \times 1 = 91$ values in each file, the same as before. Figure 2c shows the result of merging all the corresponding variables and dimensions into logical objects. The logical dimension, Months, is a composite of the virtual dimension, Month, from each file. The logical variable, TEMPERATURE, contains $7 \times 13 \times 3 = 273$ data values.

The advantage in using logical variables and dimensions is that once they are created, the user no longer has to know which files are needed to retrieve a subset of the data values for that object. The complexity of working with multiple files is reduced so that the user can think of terms of a project rather than about individual files. Using virtual dimensions gives the user the ability to increase the dimensionality of variables when dimensions are implied by the division of data into files.

3.2 Other Data Management Functions

EDM performs a large number of tasks within Envision. In addition to handling variables and dimension objects spanning multiple files, it manages the metadata associated with these objects. EDM also controls the flow of data values between the data files and visualization and analysis modules connected to it.

4. USER INTERFACE

EUI is the principal tool that a scientist uses to work with a data set in an Envision project. EUI is intended to be an easy-to-use, intuitive way to browse both metadata and data. The metadata describes data values and includes information such as collection or generation of the data and its

subsequent manipulation. Examples of simple metadata are variable name and units. More complex metadata might include a list of processing steps used to generate a particular set of data values.

4.1 User Interface Design

EUI is a client process that connects to EDM. The interface is a point-and-click environment written in X/Motif. The main part of the user interface is a table display of dimensions and variables, as shown in Figure 3. At the top of the window is the menu bar, which has a series of pull-down menus for working with projects, data files, the appearance of the table, metadata access, and visualization. In the table display, dimensions (independent variables) are represented as columns, and variables (dependent variables) are shown as rows. Raised buttons indicate which dimensions are used by which variables. Where defined, the unit attribute for each dimension and variable in the table is displayed along with its name.

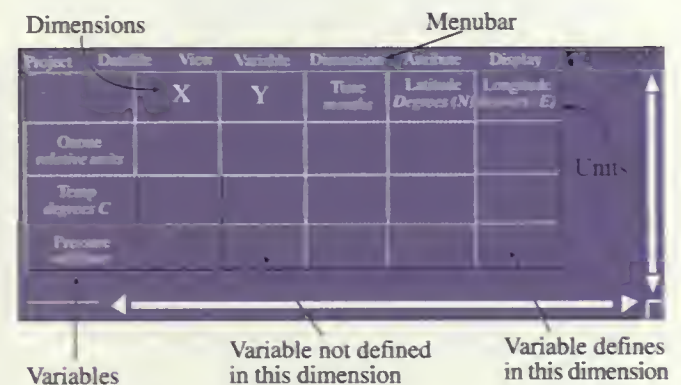


Figure 3. Envision user interface.

The table display is also customizable by the user through the creation of different views of the project. In Envision, a view is a subset of the variables in a project. Within a view, the order of rows and columns may be changed, and variables not being used may be hidden. In this way, a user can reduce the size of the table display and can efficiently browse through a large project by looking only at the variables of interest. Also, a set of views in a project can be created for different purposes. For example, if several people are working with the same project, each may wish to

see the table with different contents and arrangements.

4.2 Metadata Browsing and Editing

One of the key uses of EUI is to get access to the metadata associated with a project. The usual way to look at metadata in the past was to dump the contents of a data file using a utility such as `ncdump`. Envision makes metadata browsing more efficient through interactivity. The table display itself shows the overall structure of metadata for a project, a compilation of a set of data files. EUI also has extensive reporting features, which give pop up windows containing variable, dimension, and global attributes for the entire project or individual data files. Several levels of information detail can be switched among windows using a push button approach.

Another important feature of EUI is metadata editing. Metadata may be deleted, changed, or added to using pull-down menu options. There are a variety of applications of these capabilities. One use is the option to rename variables. For example, a variable named "T" could be changed to Temperature. If there were no units defined for a dimension or variable, these could be added. Also, if the metadata contained errors, these could be quickly corrected using EUI.

As was mentioned previously, a data set may be on a read-only medium, such as a CD-ROM. In a situation like this, it is clearly not possible to modify anything in the data files. Because Envision stores metadata in a project file, metadata on a read-only medium can still be modified or added to. Even when users are capable of modifying the data files, they may prefer to avoid the overhead of rewriting the files. A disadvantage in duplicating metadata in Envision project files is that it uses more disk space. There is also a risk that someone could change the data files outside Envision without updating the project files. To combat this possibility, Envision provides the option to write metadata changes back to the data files. A user also can create new data files from subsets of variables and their dimension ranges in a project. From a conceptual viewpoint, it is not critical where the metadata are stored, as long as they can be conveniently accessed.

5. VISUALIZATION

Envision does not perform any visualization itself, but acts as a data server for visualization and analysis tools.

5.1 How Envision Does Visualization

Creating a display in Envision involves a few steps. First, a variable is identified in EUI for further analysis. The dimensions are then selected for display axes by clicking on the raised buttons intersecting the variable row. For example, for a two-dimensional visualization, X and Y axes would be chosen. A display program, such as Collage or IDL, is chosen from a pull-down menu. At this point, another window interface pops up. This interface allows the user to restrict the ranges of the dimensions and set options for the visualization display. When the user requests a display, the data values are fetched from the data files and assembled in the proper order by EDM, then sent directly to the visualization program. Other ranges or options can be set again and new displays created, without going back to EUI.


5.2 NCSA Collage and XImage

Two public domain visualization programs used with Envision are Collage and XImage, which were written by National Center for Supercomputing Applications (NCSA). Both are raster image viewers that can display single frames or animations. XImage is several years old and has been largely replaced by Collage in the last year or so. Figure 4 shows an example of visualization options in Collage using TOMS ozone data. Figure 4a is the Collage options interface, which is written in X/Motif specifically for Envision. The other windows in the figure are from Collage. They illustrate a variety of available tools.


Another important feature of Collage is that it is a collaborative tool. This means that all the windows produced in the program can be simultaneously viewed on different networked workstations. This gives distant users or those just down the hall the ability to analyze and browse the same data sets together.




(a)




(b)



(c)



(d)



(e)

Figure 4. Visualization with Collage—(a) Collage options interface; (b) Collage interface; (c) image with annotations; (d) area histogram; (e) profile.

5.3 IDL

Another program with which Envision works is IDL, which is an interactive environment used to perform calculations on and visualizations of scientific data. IDL is used in Envision to create several visualization modules that draw one-dimensional and two-dimensional plots. Figure 5 shows an example of a two-dimensional contour display. The interface in Figure 5a is similar in function to the Collage options interface in Figure 4a, but is written entirely by IDL widget functions. The two IDL programs serve as examples to IDL users of how they can use EDM and EUI together with IDL functions to create an enhanced working environment.


6. STATUS AND PLANS

6.1 Getting Envision

Envision is a public domain software package and is available by anonymous FTP. The primary site is vista.atmos.uiuc.edu in the directory `pub/envision`, and the secondary site is csrp.tamu.edu in the directory `pub/envision`. Available at both locations are software binaries and source, documentation, and sample data sets. Instructions for obtaining NetCDF, HDF, Collage, XImage, and IDL are also provided there. For additional information, contact Keith Searight at k-searight@uiuc.edu.



(a)



(b)

Figure 5. IDL contour plot—(a) interface built with IDL widgets; (b) two-dimensional contour plot.

6.2 Current Work

The first Envision beta version was released in June 1993 for IBM RS-6000, Sun, HP, and SGI workstations. Work under way includes adding new IDL modules for three-dimensional surfaces and geographic displays with continent outlines using different map projections, as well as a port to the DEC Alpha using OSF/1. Another important feature currently being developed is a generic library for sending and receiving messages with EDM. This will allow users to easily combine their own visualization and analysis modules with Envision. Improved ways to automatically merge dimensions and variables are also under investigation.

REFERENCES

- Brown, S.A., M.J. Folk, G. Goucher, and R.K. Rew, Software for Portable Scientific Data Management, *Computers in Physics*, 7(3), 304-308, 1993.
- Rew, R.K. and G.P. Davis, The Unidata netCDF: Software for Scientific Data Access, Preprints, Sixth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, Am. Meteor. Soc., Anaheim, CA, 33-40, 1990.
- Searight, K.R., S. Pathi, K.P. Bowman, D.P. Wojtowicz, R.B. Wilhelmson, and J.E. Walsh, Envision: An Interactive System for the Management and Visualization of Large Geophysical Data Sets (abstract), *Eos Trans. AGU*, 74(16), Spring Meeting Suppl., 267, 1993a.
- Searight, K.R., D.P. Wojtowicz, K.P. Bowman, R.B. Wilhelmson, and J.E. Walsh, A Collaborative Analysis and Display System for Large Geophysical Data Sets, Preprints, Ninth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, Am. Meteor. Soc., Anaheim, CA, 72-75, 1993b.

Geophysical Phenomena Classification by Artificial Neural Networks

M.P. GOUGH and J.R. BRÜCKNER

*Space Science Centre, School of Engineering,
University of Sussex, Brighton, England*

Space science information systems involve accessing vast data bases. There is a need for an automatic process by which properties of the whole data set can be assimilated and presented to the user. Where data are in the form of spectrograms, phenomena can be detected by pattern recognition techniques. Presented are the first results obtained by applying unsupervised Artificial Neural Networks (ANNs) to the classification of magnetospheric wave spectra. The networks used here were a simple unsupervised Hamming network run on a PC and a more sophisticated CALM network run on a Sparc workstation. The ANNs were compared in their geophysical data recognition performance. CALM networks offer such qualities as fast learning, superiority in generalizing, the ability to continuously adapt to changes in the pattern set, and the possibility to modularize the network to allow the inter-relation between phenomena and data sets. This work is the first step toward an information system interface being developed at Sussex, the Whole Information System Expert (WISE). Phenomena in the data are automatically identified and provided to the user in the form of a data occurrence morphology, the Whole Information System Data Occurrence Morphology (WISDOM), along with relationships to other parameters and phenomena.

1. INTRODUCTION

Automatic classification of geophysical data is needed because of the vast quantities of data being generated and because of the nature of the data—often multidimensional and interrelated to other data sets. There are several existing and developing space data base access tools, such as the European Space Information System (Query & Correlation Environment) and the Space Physics Query Language for use with data that, in general, have already been preprocessed into simple two-dimensional form, or derived parameters.

However, much of the data generated by modern instruments are in three-dimensional form (or higher dimensions), such as wave frequency spectrograms and particle energy spectra against time. Phenomena are often difficult to extract by simple algorithms to give derived parameters, but they can generally be extracted visually by trained data analysts. For example, multiple simultaneous wave emissions are difficult to separate, as are weak signals close to the background level, such as the continuum emissions shown in the example used below.

Data analysis person-hours are usually limited by strategic reasons so that only a limited part of any given data set is ever analyzed, and with little by way of global morphologies produced. Fortunately, the recognition of phenomena in three-dimensional color spectrograms is similar to the traditional problem of image recognition for which ANNs have been increasingly used in recent years. ANNs open up the possibility of automatically analyzing the full data set.

Presented here are the first results obtained by applying two types of ANNs to the same geophysical data example. Principles of feature identifica-

tion and automatic classification are demonstrated, paving the way for the whole data set to be searched to provide a complete feature morphology. Scientific phenomena are identified as interrelated feature phenomena of different data sets.

2. GEOPHYSICAL DATA EXAMPLE USED

The data set chosen is that of natural wave emissions measured in the vicinity of the Earth by the ESA satellite GEOS-1 wave experiment consortium, S-300. The particular data example, Figure 1a, was obtained during an inbound part of the orbit of GEOS-1 where the spacecraft passes through a continuum emission generation region at 04:00 UT on day 63, 1977. The data plotted are from a narrow-band wave filter that was stepped from 0 to 77kHz in 256 steps of 300Hz, taking 22s per frequency sweep. Natural electrostatic waves increase in frequency during the plot, reflecting the increases in local plasma frequency and local electron gyrofrequency as the spacecraft moves closer to Earth into regions of higher plasma density and stronger magnetic field.

At 04:00 UT, the electrostatic wave intensity becomes so strong that there is some conversion to weak electromagnetic continuum waves, or radio waves, that propagate away from this region. Indeed, these electromagnetic continuum waves are seen well before the generation region is reached as a well-defined chevron pattern in the frequency-time spectra. This pattern results from the fact that these radio waves are propagating directly away from the source region and are modulated by the satellite-receiving antenna spin. This modulation beats against the frequency sweeping to give the observed chevron pattern. Also present in the data at high frequency is an interference line generated by the telemetry and voltage converter clocks with further continuum emissions superimposed, while at the lowest frequencies there are strong electromagnetic emissions present throughout this data example.

3. SIMPLE UNSUPERVISED HAMMING NETWORK CLASSIFICATION

The first network used is a simple unsupervised Hamming network similar to the supervised Ham-




Figure 1. Classification—(a) GEOS wave data classification example; (b) simple unsupervised Hamming network classification.

ming network flown within the SPREE instrument on STS-46, but modified for self-learning and operation on a PC. A detailed description of the original ANN is given in Gough (1993). A brief description of the modified ANN and the method of application are presented here.

Data preprocessing consisted of a simple conversion of the multibit spectra into a 1-bit image by comparison of individual data points with the local average intensity around that point on the frequency-time spectra. Preprocessed data in this example are shown later in Figure 4. The Hamming ANN consists of 16-RAM neurons each with a 10-bit-wide address corresponding to 1,024 address locations (each of 16 bits) per neuron. The total of 160 address bits were mapped with a fixed random connection to a box that was scanned across the preprocessed frequency-time spectrogram. All neuron RAM contents were initialized to zero, before learning. For each position of the box on the preprocessed data image, the image bits then defined, via the fixed mapping, individual locations within the 16 neurons. The contents of these 16 locations were read out. If the contents of a particular RAM were zero then other addresses were searched out up to a limiting Hamming distance from that address to find the closest address with non-zero contents. During learning, RAM data values were compared and a score generated for each, depending on how many of the neurons had the same value. At this stage, a value is entered into all RAMs at the locations addressed by the data covered by the box. If the score was less than 4, a previously unused negative number is the value stored. If the score was greater or equal to 4, and the value of the RAM with the highest score was negative, then a new, previously unused positive number is the value entered. If the score was between 3 and 10, and the value of the RAM with the highest score was already positive and non-zero, then that positive number is the value entered into all RAMs. The system described here has some similarities to the Sparse Distributed Memory of Kanerva (1986) and to the *n*-tuple RAM neurons of Aleksander and Stonham (1979).

In this way, features in the data initially enter the RAMs as negative numbers, but once these features recur, they are transferred to a recognized feature class with a positive number. Although the

negative values can grow to large sizes with recycling through the negative number range, the data example shown below only generated positive numbers up to the value 4 (four classes identified). Preprocessing, data scanning, and learning are fully automatic in this scheme. Once learning has finished unseen data can be classified by again scanning the box, but this time only reading RAM contents and only choosing the positive RAM value with the highest score. Although searching through RAM addresses within the limiting Hamming distance is achieved quickly, it can be eliminated by annealing the RAMs. Annealing consists of replacing all zero and negative values by the closest positive value (nearest in address Hamming distance). All of the major wave features present in Figure 1a were automatically separated by the unsupervised Hamming ANN into four different feature phenomena, denoted by color in Figure 1b. Retrospectively, an expert data analyst can easily name these four classes of feature phenomena: red—electrostatic cyclotron harmonic waves governed by the local plasma density and local magnetic field strength; blue—low-frequency electromagnetic emissions; yellow—electromagnetic continuum waves propagating away from the generation region; and green—local interference line from voltage converters/telemetry switching in the spacecraft. All the salient phenomena have been classified by this network, including separation of the weak higher frequency continuum when it is superimposed on the interference line (e.g., around 03:40 UT).

4. UNSUPERVISED CALM CLASSIFICATION

CALM was developed to address problems in ANNs, such as lack of stability, lack of speed, inability to both discriminate between and generalize over feature vectors (hereafter, vectors), and catastrophic interference [Murre, 1992]. The module is related to Grossberg's ART neural networks [Carpenter and Grossberg, 1987]. A CALM module can be described as a clustering mechanism that is able to discover statistical regularities in a stream of vectors. The module acts like a negative filter, subtracting known vectors from such a stream, so that novel vectors are

learned more intensively than the ones already known.

CALM consists of a layer of input nodes and a layer of output nodes connected to each other in a special way. During the learning phase, CALM categorizes every presented vector into a single cluster, represented by an output node. The number of clusters CALM can use for categorizing is limited by the amount of output nodes given.

To increase the stability, learning can vary from elaboration learning (about 200 learning cycles for ANN conversion) to activation learning (about 50 cycles). A novel vector causes elaboration learning with the formation of new associations, whereas a known vector just strengthens the existing associations. A high learning speed is achieved by using arousal effects, node competition, and noise within a module. The last ensures that convergence is reached with every vector, as it resolves competition deadlocks in the module. The combination of single modules with a hierarchical structured ANN eliminates the effect of catastrophic interference if none of the module's capacity is exceeded. This means that well-learned vectors are not replaced when new vectors are learned and the old vectors are not presented again.

The last section has shown how a crude classification can be achieved with simple preprocessing effort and a simple ANN. Current work concentrates on improving the preprocessing techniques to obtain abstract vectors for a classification with the more sophisticated CALM. The following example will show how abstract vectors with intensity, smallest frequency, orientation, and circularity components are gained from the example data set. The categorization result for a module with 7 input and 10 output nodes, to accommodate 7 vector components and a maximum of 10 clusters, will then be presented.

Preprocessing the spectrogram to the features desired is a sequential process. An optional smoothing of the spectrogram by low-pass filtering precedes a two-step segmentation. First, local background dependent filtering, such as in the last section, yields a black-and-white picture where black areas denote local maxima. These areas represent features but cannot be used solely to extract all vector components needed. Second, sobel edge detection and conditional bit-wise copying of pixels from the original spectrogram to

the black-and-white picture result in a picture containing all features surrounded with a black boundary against a white background, shown in Figure 2a. The feature areas are then identified by a simple search algorithm, which allows a flood fill-related algorithm to find the pixels within the boundary and the boundary coordinates for each area. While the algorithm fills an area, it evaluates a pixel intensity histogram and stores the boundary coordinates. Thus, the vector component's intensity, smallest frequency, orientation, and circularity can easily be computed and normalized for all features to yield a set of vectors. Here, three intensity components were gained from the intensity histogram and the smallest frequency component from the minimum boundary Y-coordinate. The feature angle from 0 to 90 degrees to the horizontal was used for two orientation components, and the circularity component was the ratio of the number of pixels inside the feature area to the square of the number of boundary pixels.

Unfortunately, the vector set cannot be presented to CALM right away. The components need to be weighted (multiplied by a parameter) to achieve an optimal classification result. Suitable parameters are found by using a simple genetic algorithm [Goldberg, 1989] with a human classification as reference. A genetic algorithm is a mechanism that uses random choice as a tool to guide a highly exploitative search through a coding of a parameter space. With such optimized vectors, an untrained CALM found clusters as shown in Figure 2c. Different colors represent different clusters and hence phenomena classes, as for section 3. The result is quite similar to the ideal classification of a data analyst given in Figure 2b.

As the learned information is contained in the network's learning weights, it is possible to store and restore these (pretrained network). This permits a quick scan through further data sets with or without additional elaboration learning. Both identify new, previously unknown features, where the former classifies all features and learns new features (continuous learning) and the latter quickly classifies all features (within 3 seconds for this data set). Apart from defining what kinds of vector components to use and an initial, genetic algorithm-supported, component parameter setting, the above processes are fully automatic.

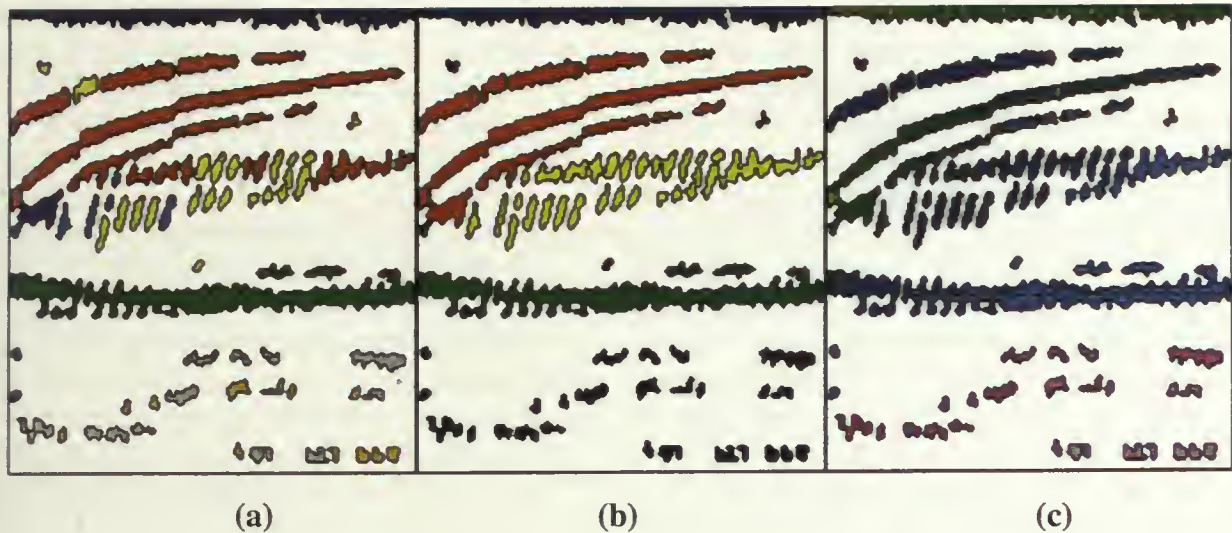


Figure 2. CALM example—(a) preprocessed spectrogram; (b) data analyst classification; (c) ANN classification. The data analyst has classified the features from the preprocessed spectrogram into four different classes with the colors used in Figure 1(b). Black features denote unclassified features. The ANN classification shows how four known and five new classes (new colors) were found by CALM. Note that different color scales were applied in Figure 1(a) and in the preprocessed spectrogram here.

There are two false classifications in the red class. The small blob has been mistaken as a yellow and the thick line as a blue member from similar intensity distributions, angles, and circularities. In these cases, the Euclidean distance to members of the false class was smaller than to members of the correct class or too small to meet the discrimination threshold (see below). CALM has separated the yellow class into yellow and turquoise subclasses. Both subclasses differ slightly in their intensity distribution. Some of the yellow members were classified as red members because of inadequate segmentation, which resulted in false angle components.

A straightforward definition of classification accuracy, where a human classification is taken as a reference, shown in Figure 2b, and the percentage of correctly categorized features for each class is computed, shows that 62 percent of the yellow, 100 percent of the blue, 84 percent of the red, and 100 percent of the green class features were correctly classified. The overall accuracy is 87 percent.

5. COMPARISON OF ANNs USED

The relatively crude Hamming network used first has the advantage of simplicity and reliable

effective classification. It is quite fast, taking 1 minute on a 33MHz 80486 to learn the data example of Figure 1, and has low memory requirements, with the whole PC demonstrator program taking less than 1 Mbyte. However, with the 1-bit preprocessed image, it is only as good as the preprocessing algorithm used, but an effective algorithm can be found for most space plasma physics data types. The only user input is to select the size of the scanning box and the algorithm for pre-processing. In both cases, these selections, once optimized by trial, would be valid for all data of that data type in the data base. Major disadvantages of this simple ANN are lack of scaling to data feature size in the image and lack of sensitivity to intensity differences between feature types.

The developed software package for preprocessing and CALM was run on a Solbourne S4000 with a Sparc 1 processor. The preprocessing differs from the one in the last section. Here, each feature is represented by seven scaling independent, normalized, real-value components. Thus, the main feature properties are emphasized. There is no limit of dimensionality, so that one can use as many descriptors (components) as needed. The above result suggests the use of more and/or better descriptors to avoid misclassification. The discrimina-

tion threshold of 0.8 for the minimum Euclidean distance to categorize vectors in different clusters, along with a maximum component value of 1.0, shows why this might be necessary. The preprocessing can be as fast as the machine and an optimized program code allows; here, it was around 1.5 minutes for the example spectrogram. So far, little effort has been applied to code optimization. Categorizing of the vector set, which contained 77 vectors, took 14 seconds. A smaller number of output nodes than clusters causes the network to generalize over the input vectors. This may be an advantage if a coarse classification is required or a disadvantage if fine classification is needed. In this example, the number of output nodes was larger than required, which allowed the categorization of new features. As CALM modules can be combined to increasingly complex hierarchical ANNs, it is possible to find more abstract classes within interrelated phenomena. If the classes for a first layer of modules were *letters* (features), a second layer could combine them to *words* (feature phenomena) and a third layer to *sentences* (scientific phenom-

ena). It might be necessary to tune the dynamics of individual modules in each layer to yield an optimal result as, for example, the generalization behavior also depends on the module parameters. It has been shown that genetic algorithms can help here as well [Murre, 1992].

6. PROPOSED SYSTEM: WHOLE INFORMATION SYSTEM EXPERT (WISE)

Repeated application of ANNs to data bases can provide a degree of automatic data analysis. Work at the University of Sussex is aiming toward a system known as WISE [Gough, 1992]. In this system, illustrated in Figure 3, ANNs are applied to data sets in a series of stages:

Stage 0: The user (expert) optimizes the preprocessing needed for any given data set before the use of ANN. In general, this is conversion of data into image form without loss of essential features. Feature vectors are then generated.

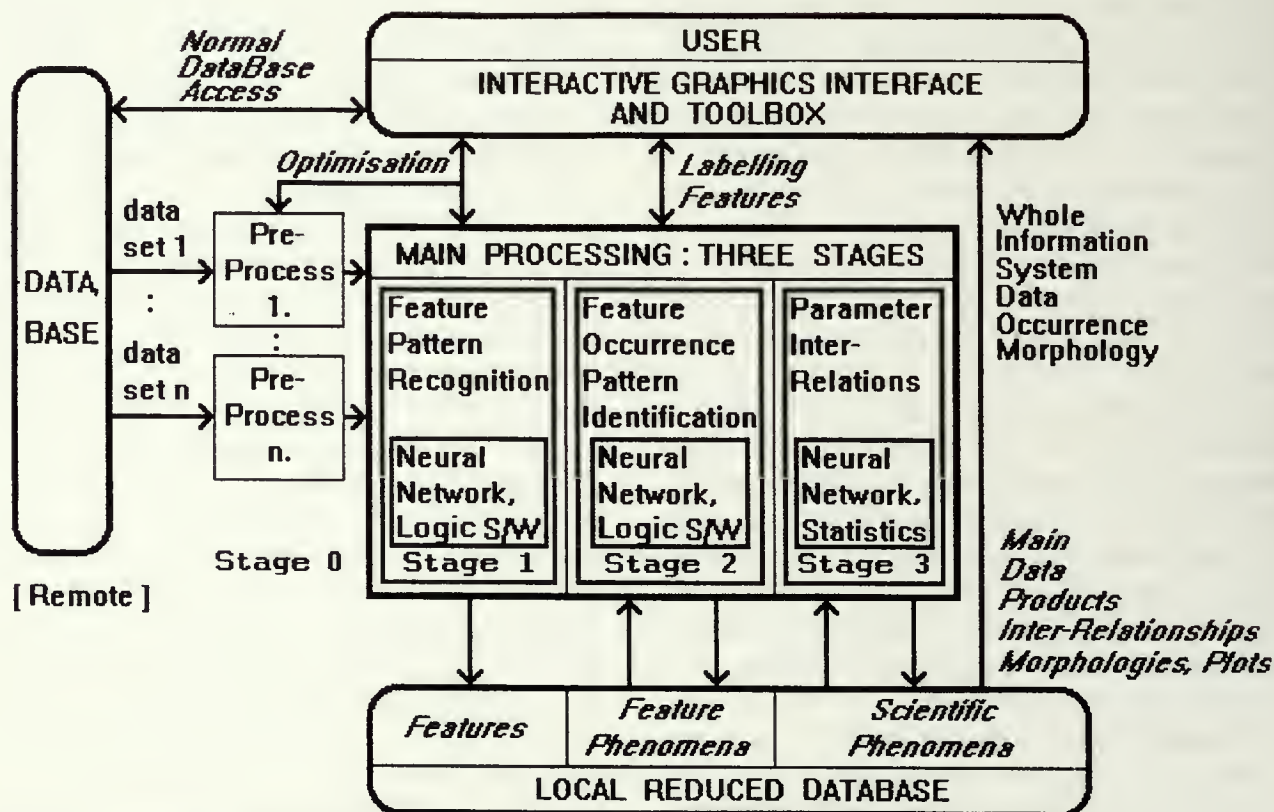


Figure 3. Schematic diagram of WISE system.

- Stage 1: Unsupervised ANN is applied to the feature vectors to classify features present in the image, similar to the traditional use of ANN in image/pattern recognition. A local data base stores ANN weights, and thus features are recognized.
- Stage 2: The patterns of occurrence of features in the data set are themselves classified by ANNs, statistics packages, and/or logic programming in a similar way, and again stored in the local data base as feature phenomena.
- Stage 3: Similar methods as in stage 2 are used to find interrelationships between feature phenomena in the same or different data sets to identify scientific phenomena. The final results are then displayed, for example, like GEOS morphologies (see section 8).

WISE would include toolkits for signal and image processing, neural networks, and statistics packages to achieve the above stages. The main product is the Whole Information System Data Occurrence Morphology (or WISDOM). An expected scenario of WISE use is as follows. The expert familiar with each data set, probably the instrument PI or CoI, optimizes preprocessing and runs WISE unsupervised on the whole data set. This expert labels features found and related feature groups, or phenomena, with the names known in that community of researchers. It is expected that such a system could reveal new phenomena or relationships at this stage. Later, a nonexpert user could apply the trained WISE system to identify any feature in the data set and instantly be given its morphology, or importance, along with all related features. In this way, space-acquired data bases can be better utilized by nonexperts and, being opened up to a wider user community, would become more fully utilized. At present, it is often estimated that less than 10 percent of all space plasma physics data has been accessed.

7. INITIAL WISE FRONT-END PC DEMONSTRATOR

A prototype front-end for WISE was written incorporating the simple unsupervised Hamming

network running under Windows on a PC and demonstrated at AGU in the spring of 1993. It can be seen from the main screen in Figure 4 that the preprocessing and ANN data classification can be fully monitored as processing occurs, with learning statistics keeping the user updated on patterns that are learned. At any time, the mouse can be used to test the response at a given location in the data. It is also possible to manually train the network by way of chosen examples (supervised learning). In this case, it was found that considerable classification was possible after showing only one example of each of the four main data classes present in the data of Figure 1. Other options include the selection of further windows to provide displays of the typical inputs corresponding to a given class, various classification statistics summaries, and displays of information distribution within the neurons with the possibility to anneal or trim the network for enhanced performance and minimized class overlap.

8. TYPICAL WISDOM PRODUCTS

In this concluding section, it is shown how such a system can be used to give phenomena morphology across the whole data set, and hence, an understanding of the relative importance of any given phenomenon. Once an unsupervised ANN has identified a series of phenomena classes in a data set, it is a simple matter to then process the whole data base to generate morphologies of phenomena occurrence across the data base.

Figure 5 shows an example of possible use of the future WISE system to understand the morphology of GEOS electrostatic emissions near the plasma frequency. Note that in these plots emissions have been extracted by a dedicated numerical algorithm [Gough et al., 1980], but, from the above discussion, ANN would perform with equal satisfaction. In the examples shown in Figures 5a and 5b, 430 days of GEOS II wave data have been analyzed. Data are plotted in polar plots against magnetic local time (midday at top, dawn to right, midnight at bottom, and dusk to left).

Plots such as these immediately show the data analyst where wave emissions are strongest in geomagnetic coordinates, how the power varies with geomagnetic activity, and how plasma fre-

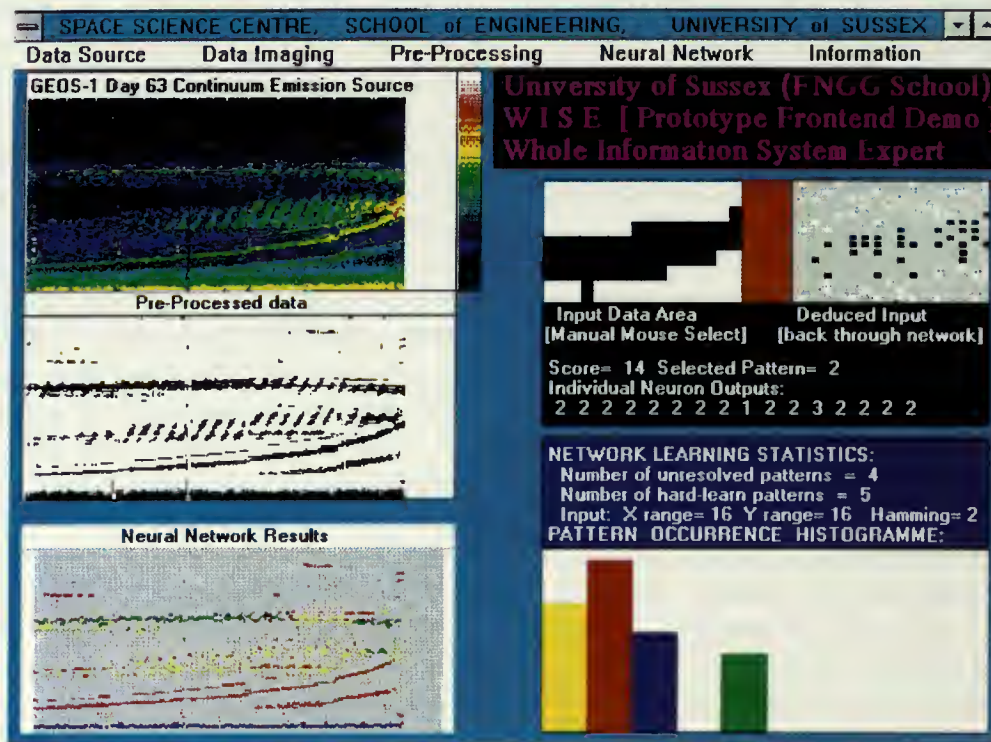


Figure 4. PC screen showing use of WISE front-end demonstrator. On the left are shown top to bottom: original wave spectrogram, 1-bit preprocessed data, and the results from applying ANN. Clicking the mouse on any position (on any of these three windows) provides in the top right window a closeup view of the selected area, with neuron scoring and subsequent feature class identification. The lower right window provides a histogram of occurrence of the feature classes in the data. Pull-down menus provide full control over data access, preprocessing, and neural network configuration.




Figure 5. Electrostatic wave morphology in the GEOS II data set—(a) GEOS II wave power plotted radially against LT; (b) GEOS II wave frequency plotted radially against LT.

quency and hence cold electron density vary with local time and geomagnetic activity. An equivalent study by a human data analyst would require many intensive person-years of effort because of the need to accurately identify wave phenomena in such a study.

REFERENCES

- Aleksander, I. and T.J. Stonham, Guide to Pattern Recognition Using Random-Access Memories, *IEEE J. Computers and Digital Techniques*, 2, 29, 1979.
- Carpenter, G.A. and S. Grossberg, ART 2: Self-Organization of Stable Category Recognition Codes for Analog Input Patterns, *Applied Optics*, 26, 23, 1987.
- Goldberg, D.E., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, 1989.
- Gough, M.P., P.J. Christiansen, and E. Gershuny, Electrostatic Wave Morphology Near the Geostationary Orbit, in *Space Research, Advances in Space Exploration I*, edited by K. Knott, Pergamon, 1980.
- Gough, M.P., Automatic Feature Recognition and Morphology Generation, paper presented at Earth and Space Science Information Systems, Pasadena, February 1992.
- Gough, M.P., Space Instrument Neural Network for Real-Time Data Analysis, *IEEE Trans. Geoscience and Remote Sensing*, 31, 6, November 1993.
- Kanerva, P., Parallel Structure in Human and Computer Memory, in *Neural Networks for Computing*, edited by J.S. Denker, 247-258, American Institute of Physics, New York, 1986.
- Murre, J.M.J., R.H. Phaf, and G. Wolters, CALM: Categorising and Learning Module, *Neural Networks*, 5, 55-82, 1992.
- Murre, J.M.J., *Learning and Categorization in Modular Neural Networks*, Lawrence Erlbaum, 1992.

DITDOS: A Set of Design Specifications for Distributed Data Inventories

T.A. KING, R.J. WALKER, S.P. JOY

Institute of Geophysics and Planetary Physics, University of California, Los Angeles, CA

The analysis of space science data often requires researchers to work with many different types of data. For instance, correlative analysis can require data from multiple instruments on a single spacecraft, multiple spacecraft, and ground-based data. Typically, data from each source are available in a different format and have been written on a different type of computer, and so much effort must be spent to read the data and convert it to the computer and format that the researchers use in their analysis. The large and ever-growing amount of data and the large investment by the scientific community in software that require a specific data format make using standard data formats impractical.

A format-independent approach to accessing and analyzing disparate data is key to being able to deliver data to a diverse community in a timely fashion. The system in use at the Planetary Plasma Interactions (PPI) node of the NASA Planetary Data System (PDS) is based on the object-oriented Distributed Inventory Tracking and Data Ordering Specification (DITDOS), which describes data inventories in a storage independent way. The specifications have been designed to make it possible to build DITDOS compliant inventories that can exist on portable media such as CD-ROMs. The portable media can be moved within a system, or from system to system, and still be used without modification. Several applications have been developed to work with DITDOS compliant data holdings. One is a windows-based client/server application, which helps guide the user in the selection of data. A user can select a data base, then a data set, then a specific data file, and then either order the data and receive it immediately if it is online or request that it be brought online if it is not. A user can also view data by any of the supported methods. DITDOS makes it possible to use already existing applications for data-specific actions, and this is done whenever possible. Another application is a stand-alone tool to assist in the extraction of data from portable media, such as CD-ROMs. In addition to the applications, there is a set of libraries that can facilitate building new DITDOS compliant applications.

1. INTRODUCTION

Space physics researchers often are confronted with data in a variety of formats. Many information system designers also are faced with this problem, in both the commercial and scientific communities. In the scientific community, there are a number of reasons why data obtained from different sources use different storage methods. In general, projects adopt a formal storage format, but these formats differ from project to project. In addition, the binary representation of the data can differ among projects. These differences occur because different projects use different computing environments and design their formats to address specific needs. For older missions, the computing environment in the scientific community may have changed so much that it is difficult to transform the original data into a usable form.

There have been attempts in the past to develop universal formats or to impose a fixed computing environment. These efforts have encountered limited success, primarily because of the cost of imposing standards retroactively. Restoring, reformatting, and re-archiving of a historical project's data holdings can be a substantial fraction of the dollars available for research on the data. Some balance between making data readily available and easy to use and providing research support must be reached. The Planetary Plasma Interactions (PPI) node of the NASA Planetary Data System (PDS) has been striving to provide easy access to data from all the fields and particle instruments on planetary missions while preserving the investments made in software analysis and research tools by individual projects, as well as providing a pathway to new technologies and new applications. We are doing this by adopting a format-and-operating

system-independent approach to accessing and analyzing data.

2. MODEL DATA SYSTEM

In 1988, when we designed the first version of the PDS/PPI node software, we adopted a model for a distributed data system that consisted of multiple access points to (mostly) centralized data holdings. This model was adopted primarily to address slow user response times typical when systems with centralized access and data holdings were accessed remotely. In addition, all data were reformatted into a single, common format when it was placed within the system. This was effective, but we quickly learned that there were important issues this model did not address. First, because we used a standard data format that was different from the one used by the project, when members of the project team obtained the data, they could not read the data with their existing software. This was very inconvenient. To address this we needed to provide some translation capability from the standard format to the specific format used by the project. Second, we found we were spending significant resources converting data from a specific format to the standard format. These resources included personnel to perform the conversion, space to store the converted copies of the data, and the additional management overhead of tracking various versions

of the same data. Third, by maintaining a central data holding, we were not making the most effective use of pre-existing resources. We have found that by maintaining data holdings within active research institutions, the quality of the data is increased because the data are being used in research and problems are corrected as they are found. Fourth, while we had an effective method for managing our data holdings, researchers still had to use different methods to manage the data received from us. In some cases, users wanted large quantities of data for statistical studies. Our existing system provided little support for such endeavors.

Recently, we used the lessons learned from the first version of the PDS/PPI node software to design our second generation data system. It is a highly distributed data system with distributed data holdings and distributed data access points. In this system, we decided that the data holdings would remain in the data format the project or investigation deemed most suitable. Because we must be able to handle data in any format, it was not possible to include the data within a particular data base management system. We needed a standardized inventory scheme so that individual data files could be tracked.

The overall model for a system designed to address these issues is depicted in Figure 1. In the new system model, there are two kinds of servers, primary and secondary, and two kinds of clients,

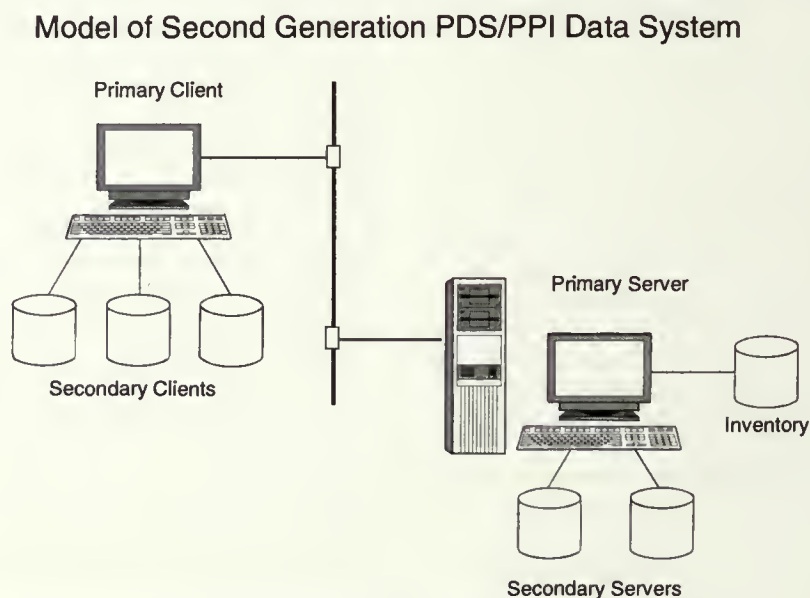


Figure 1. The data system model used in the design of the second generation PDS/PPI system.

primary and secondary. The primary server provides two functions. First, it provides standardized access to data inventories to facilitate locating the data holding of interest. In this function, the primary server delivers inventory information in a specific format and with a content described by the inventory contents specifications (to be defined later) to the requesting client. There are no restrictions on how the inventory information is physically stored. It can be stored in data base management systems or in alternative inventory management systems. The second function of the primary server is to provide a standard method for accessing data holdings. The primary server calls upon a secondary server (or server agent), which can read data in a specific format and deliver a stream of bytes to the primary server. The secondary servers are format-specific servers that deliver requested data to the primary server for transport to a requesting client. The primary server takes the stream of bytes produced by the secondary server, packetizes it, and delivers it to the requesting client. The movement of data from the secondary server to the primary server and from the primary server to the requesting client is accomplished by using specifically designed protocols.

Like the servers, there are both primary clients and secondary clients. A primary client interacts

with a primary server by using the primary client/server protocol. The role of the primary client is to provide a view of the inventories provided by the primary server and to enable the selection of a specific data holding. When a primary client is handling a data holding request, the data are delivered as a packetized stream of bytes, which are preceded by information detailing the specific data format the stream represents. This data format information is used to select a particular secondary client (or client agent) to handle the data stream. The primary client's function is limited to unpacketizing the data and delivering it to the secondary client.

Like their secondary server counterparts, secondary clients operate on data of a specific format. Secondary clients are typically identified as either a writer or a viewer, depending on the destination of the data. A writer typically writes the data to a local disk, and a viewer typically displays the data on the screen. In some cases, a viewer may also serve as a writer—for example, when a viewer allows the user to examine or select subsets of the data before writing the data to disk. Figure 2 shows the information pathways between the primary client and server and between secondary clients and servers.

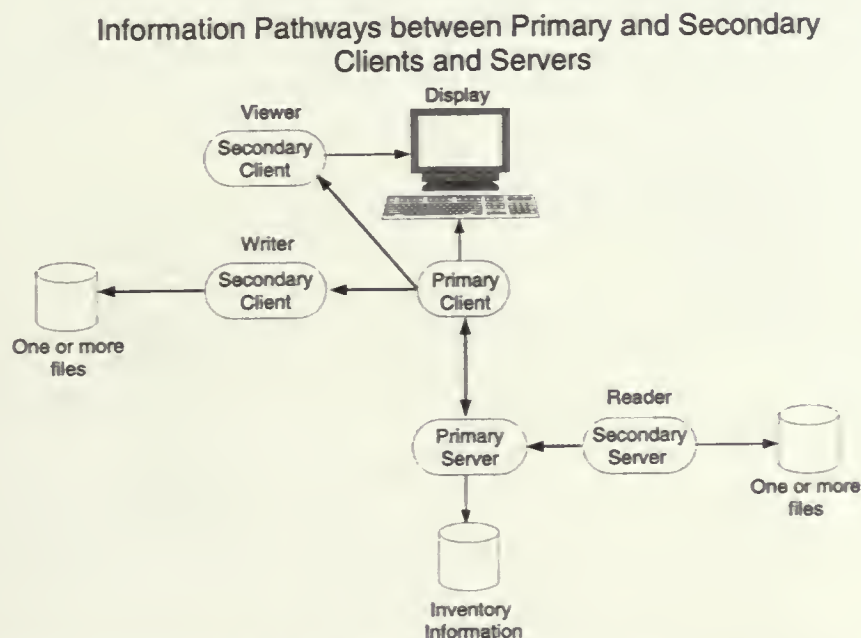


Figure 2. Information pathways between the primary client and server and secondary clients and servers. The arrows indicate the flow direction of information (data and inventory).

The entire system must be portable and scalable [King et al., 1993b]. We wanted to be able to include with large data collections a self-contained data system (portability) with which a user could manage data and that is independent of a computer and operating system. In addition, we wanted to be able to migrate collections of data online and offline with a minimal impact on the entire system (scalability) [King et al., 1993a].

3. DITDOSDETAILS

We have defined a set of specifications called the Distributed Inventory Tracking and Data Ordering Specification (DITDOS), which detail the six types of metadata required to manage data using our approach. The inventory contents are client/server protocol, secondary server protocol, primary server protocol, secondary client/server configuration, and inventory description. Each of these are discussed separately.

3.1 Inventory Contents Specification

The inventory contents specification defines the data returned by a primary server when an inventory request is made by a client. The two types of inventory information are data base and data set information. A data base description describes collections of data sets. A data set describes collec-

tions of data holdings and includes information about access privileges and curator information. Each data holding is referred to as a member of a data set. Because membership in a data set is maintained through reference to a physical data holding, it is possible for a data holding to be a member of multiple data sets without reproducing the data.

While the inventory information has a fixed format and structure and was designed to be used in a relational data base system, there is no requirement that the inventory information actually be stored in a relational data base system. The only requirement is that a primary server deliver the inventory information in a specific format. Figure 3 shows the contents of the various relational "tables," called the inventory tables, and how the tables interrelate. The contents of the inventory tables detail which data sets are available and which members exist in each data set. Each set of inventory tables is referred to as a data base. There are four "tables" in a data base: (1) the "data set" table, which defines which data sets are in the data base; (2) the "member" table, which defines which data files are members of a specific data set; (3) the "curator" table, which provides information on who contributed the data set; and (4) the "privs" table, which describes privileges users have on data sets. The detailed structure of each inventory table follows.

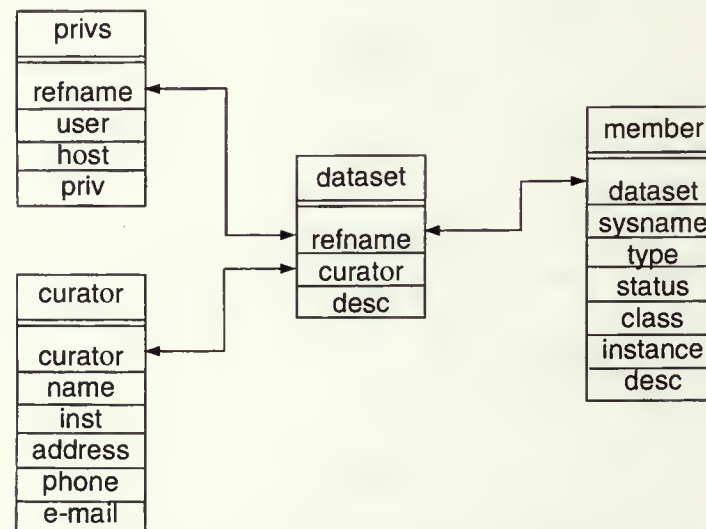


Figure 3. The required information in each inventory table and the relation of fields in each table. The lines show dependencies.

3.1.1 Data Set Table

The data set table defines which data sets exist, the curator of each individual data set, and a description of the data set contents. The structure of the data set information is shown in Table 1.

3.1.2 Member Table

The member table defines the members of each data set. It classifies the content of the data holding and identifies the storage format. It also contains a description of the contents of the data holding. The structure of the member table is shown in Table 2.

3.1.3 Curator Table

The curator table provides detailed information for each curator reference name. This includes the physical mailing address, telephone number, and E-mail address. The structure of the curator table is shown in Table 3.

3.1.4 Privs Table

The privs table identifies which users from which host have specific privileges on a data set. These privileges apply to all possible operations that can be performed on all member data holdings. The structure of the privs table is shown in Table 4.

Table 1. Data Set Table.

Field Name	Size	Type	Description
refname	64	AlphaNumeric	The external reference name for the data set.
curator	32	AlphaNumeric	The reference name of the curator for the data set.
desc	512	AlphaNumeric	A textual description of the contents of the data set.

Table 2. Member Table.

Field Name	Size	Type	Description
data set	64	AlphaNumeric	The external reference name of the data set of which this file is a member.
sysname	256	AlphaNumeric	The system name (path) to the data holding to which this entry refers.
type	32	AlphaNumeric	The data content type of a member. This is typically a grouping, such as data, document, image, animation, etc.
status	64	AlphaNumeric	A description of the status of the member file. Words such as “online” and “offline” should appear in the status description.
class	32	AlphaNumeric	A phrase identifying the class of the member file. The class is the name of the storage format.
instance	64	AlphaNumeric	A phrase identifying the instance of the member file. The instance is equivalent to the version of a format.
desc	512	AlphaNumeric	A textual description of the contents of the member.

Table 3. Curator Table.

Field Name	Size	Type	Description
curator	32	AlphaNumeric	The reference name of the curator for the data set.
name	64	AlphaNumeric	The full name of the curator.
inst	64	AlphaNumeric	Institution or affiliation name.
address	64	AlphaNumeric	The full mailing address—for example, street, city, state, zip, and country.
phone	32	AlphaNumeric	Telephone number for contacting the curator.
e-mail	128	AlphaNumeric	Full E-mail address for contacting the curator.

Table 4. Privs Table.

Field Name	Size	Type	Description
refname	64	AlphaNumeric	The external reference name for the data set for which this set of privileges applies.
user	32	AlphaNumeric	The name of the user granted this set of privileges. The entry may contain wild cards.
host	32	AlphaNumeric	The name of the host from which the user must originate for this set of privileges to be granted. The entry may contain wild cards.
priv	4	Integer	A bit map of privileges that are granted. The following fields are defined: bit 1: select; bit 2: insert; bit 3: update; and bit 4: delete. All other bit fields are undefined.

3.1.5 Notes on the Tables

With the exception of the “sysname” field in the member table, all information in a data base is an abstraction—that is, it is something that is defined to help organize and manage the underlying data. For example, a data set name (refname in the data set table) can contain any sequence of characters. It can be a serial number or some name derived by using a particular nomenclature. Unique aspects of the inventory data are the “class” and “instance” entries in the member table. This information is used to determine the specific format of the data file and in turn is used to determine which secondary client or secondary server to call to manipulate the data. For example, if the class for a member was “UCLA/IGPP flat file” and the instance was “Version 2.0,” then the format of the member file would be the one developed at UCLA recently in which the member file consisted of two physical files, one a PDS compliant label file and the other the data file.

Another unique aspect of the inventory tables involves the use of the “user” and “host” fields in the privs table. This table was designed to work easily with networks so the user and host fields support the use of wild cards. This allows privileges to be granted based on the user’s host, or the user name or some combination of user and host. For example, if the user field was “%” (% is the wild card character)

and host was “kingsun,” then any user from “kingsun” would be granted the associated privilege. Possible privileges that may be granted are: SELECT, INSERT, UPDATE, and DELETE.

All information provided in the inventory is considered opaque (read-only) to the client and is used solely for selection purposes. Information maintained in the privs table is used internally by the server and typically should not be a concern of a client. It is expected that only special clients, referred to as administration tools, would actually alter the content of the inventory files.

3.2 Primary Client/Server Protocol Specification

The primary client/server protocol is a bi-directional protocol, which uses tagged packets. A packet consists of a tag accompanied by data. The tag associated with the packet is either a request or a response, and the data associated with the packet is either additional information needed to complete the request or the results of a request when the packet is a response. Possible responses from a server are shown in Table 5.

A server may be configured by a client by sending a packet tagged at SET. The data portion of a SET packet must contain a string with a “keyword=value” syntax. The current parameters that can be set within a server are shown in Table 6.

Table 5. Primary Client/Server Protocol Response.

Tag	Description	Data
OK	Indicates that the request has been received and will be processed.	None.
ERR	Indicates that an error has occurred while attempting to act on the most recent request.	Text containing a descriptive message about the cause of the error.
DATA	Signifies a packet that contains data, which is in response to the most recent request.	The data requested.
OBJECT	Signifies a packet that contains information about the subsequent data packets. An object packet must precede all data transfers.	The record from the member table corresponding to the member.
END	Indicates the end of the data transfer. The request has been fulfilled. No more data are to follow.	None.

Table 6. Current Parameters.

Parameter	Description
DATA_TYPE	Defines the type of data of interest. Defining this parameter limits searches of the member table to members of the specified type.
SEARCH_METHOD	Defines the search method to use when locating records. Possible values are SERIAL and GLOBAL. A SERIAL search proceeds through table records sequentially. A GLOBAL search proceeds through tables by using a relational model, matching all possible occurrences.
MEMBER_SYSNAME	Defines the system name of the desired member file. This is required to eliminate ambiguities when there are multiple member files with the same data type.
TIME_FORMAT	Defines the time format that secondary servers use when returning time values. The value assigned to the variable is the “name” of the desired format. Currently, 17 different formats are defined. See text for details.

Supported time formats are:

- AMERDATE (1/19/93)
- EURODATE (19.1.93)
- ABBREAMER (Jan 19, 1993)
- ABBREURO (19 Jan 1993)
- LONGAMER (January 19, 1993)
- LONGEURO (19 January 1993)
- NUMERICAL (93.019)
- DAYNUMBER (93 019)
- JAPANDATE (93.19.1)
- NIPPONDATE (93.1.19)
- HIGHLOW (93 1 19)
- 1SEEDATE (93 19 1 19)
- DFS_STYLE (1993-Jan-19)
- ABBRDFS_STYLE (1993/01/19)
- PDS_STYLE (19930119T12:43:36.002Z)
- ISO (19930119T124336Z)
- BINARY (853677816.0002)

Each date, with the exception of PDS_STYLE, ISO, and BINARY, are followed by a time of day in the format HH:MM.SS.mmm (12:43:36.002).

Possible requests from a client are shown in Table 7.

The client/server protocol stipulates that when a client initiates a connection to a server, it must register with the server the name of the user running the client and the name of the originating host. The protocol also stipulates that only a client can initiate a request for information. A server must always acknowledge the receipt of all client requests. A server's acknowledgment can be either "OK" if the request was accepted and will be processed or "ERR" if an error has occurred and the request will not be processed. An additional requirement is that a data base must be opened before any queries can be submitted.

3.3 Secondary Server Protocol Specification

When a secondary server is called by the primary server, it is called with a single command line argument. That argument is the value of the "sysname" field in the member table for the requested member. In addition, two environment variables are set. These are:

- DITDOS_TIME_FORMAT, which is set to a value defined by the client using the "SET" packet.
- DITDOS_QUERY, which is set to an SQL query submitted with the QUERY packet.

An application is not required to use any environment variables, which are defined so that an application may provide additional services beyond data delivery.

A secondary server accesses the data file associated with the passed "sysname," converts the data file into a data stream, and writes the resulting data stream to "standard out." Because the interaction between the primary server and the secondary server is local (i.e., on the same machine), the definition of "standard out" is irrelevant. Because all access to data by a client must go through the primary server, which packetizes data in a standardized manner, all platform dependencies are transparent to the client.

3.4 Secondary Client Protocol Specification

The secondary client protocol is identical to the secondary server protocol with the following exception. A secondary client reads data from standard in and outputs data to some device (i.e., the screen, a disk file, or another process). The

Table 7. Possible Client Requests.

Tag	Description	Data
REGISTER	Signifies a packet which transfers user and host registration information.	The user and host name for the client in the format user@host.
QUERY	Requests specific information.	The query for information in SQL format.
OPEN_DB	Requests opening a data base.	The name of the database.
CLOSE_DB	Requests closing the previously open data base.	None.
VERSION	Requests the version number of the server.	None.
RESET	Requests resetting the server to an initial state.	None.
DONE	Instructs the server that the client has completed all transactions and is exiting.	None.
CANCEL	Requests cancelling any currently running queries.	None.
SET	Sets specific parameters in a server.	Parameter setting instructions in the form "parameter=value."

definition of “standard in” is platform specific, but is transparent to the server because all data must pass through the primary client on their way to the secondary client.

3.5 Secondary Client/Server Configuration Specification

Each primary client and server can launch a secondary client or server through a simple function. This function takes the class and instance information, along with the type of client or server agent required. Possible choices for agents are readers, writers, viewers, and requesters. A reader is a secondary server agent, writers and viewers are secondary client agents, and a requester is a special type of agent that can be used to automate some independent action, such as bringing offline data online. In this case, a requester could send E-mail to a specific individual or could run a process that will retrieve the requested file.

Agent categories are divided into specific types. For example, the user may have a graphic viewer and a textual viewer. The available agents for a particular class and instance are specified in a configuration file, which is external to the primary client and server. The format of entries in the configuration file is:

```
class={ class name }
instance={ instance name }
{ agent }. { type } = { application }
```

All { agent }. { type } definitions are associated with the preceding class and instance definition. If an { application } does not have a preceding path, then the application is expected to be located in the installation directory. An example of an agent configuration for the UCLA/IGPP flat file is as follows:

```
class = UCLA/IGPP FLAT FILE
instance = Version 1.0
    reader.native = ffconvert
    writer.native = byte-write
    writer.ascii = fwrite-ascii
    viewer.text = xffviewer
    viewer.graphic = xffgraph
    request.offline = /usr/ucb/mail
```

This configuration defines one reader, two writers, two viewers, and a way to request offline data. The reader is a “native” reader, which will read the flat file in its native format. The two writers are “native” and “ascii.” The native writer will write a flat file in its native format, and the “ascii” writer will convert the native flat file format to an ASCII format. The viewers allow visualizing the data in either a graphic or textual format.

A typical client application interacts with the agent configuration file to present choices to the user when specific actions are to be taken.

3.6 Inventory Descriptions Specification

Because there are occasions when a user may want to transfer a description of a DITDOS inventory in a portable way, an inventory description language was defined. An inventory description is an ASCII file containing directives that list member files and link them to data sets. The basic syntax for the inventory description is:

```
define { object } { object option... },
```

where { object } can be “data set” or “member” and { object option } is an { object } specific list of options. There also are commands for creating data bases and for setting the focus for subsequent define commands.

3.7 Environment Variables

Currently, only one execution time environment variable is defined. The name of the variable is “DITDOS” and must be assigned the path to the installation directory. Client applications use the environment variable “DITDOS_OPERATOR” to determine whom to send comments.

4. CONCLUSIONS

The DITDOS specifications provide a protocol for designing portable applications for accessing data inventories. The non-intrusive characteristic of systems based on DITDOS makes it easy to add distributed access to preexisting data systems. In addition, the scalability of DITDOS-based systems makes it possible to build self-contained data

collections that include DITDOS inventories. These self-contained data collections can be put on hard disk, CD-ROM, or other distributable media. In addition, a single method of access can be used to obtain data from online data holdings or from CD-ROMs. Various applications have been written based on the DITDOS specifications and are currently in use as the data access system for the PPI/PDS node. These include a server for inventories maintained in UCLA/IGPP flat files, an inventory description translator, an X-Windows client (which provides for data viewing and ordering), and command line extraction tools to accompany CD-ROMs.

Acknowledgments. Funding for this work was provided by the NASA PDS through contract JPL959026. This is IGPP Publication 4097.

REFERENCES

- King, T., S. Joy, and R. Walker, The Design and Implementation of Scalable Data Systems and Incremental Data Sets, in *Proceedings of the International Science Year Conference on Earth and Space Science Information Systems*, edited by A. Zygielbaum, 461-466, AIP Press, Woodbury, New York, 1993a.
- King, T., S. Joy, and R. Walker, The Morphology and Architecture of a Distributed Data System, in *Proceedings of the International Science Year Conference on Earth and Space Science Information Systems*, edited by A. Zygielbaum, 454-460, AIP Press, Woodbury, New York, 1993b.

Educational Software for the Visualization of Space Plasma Processes

C.T. RUSSELL, G. LE, J.G. LUHMANN, and B. LITTLEFIELD

*Institute of Geophysics and Planetary Physics,
University of California, Los Angeles, CA*

The UCLA Space Physics Group has developed educational software composed of a series of modules to assist students with understanding basic concepts of space plasmas and charged particle motion. Present modules cover planetary magnetospheres, charged particle motion, cold plasma waves, collisionless shock waves, and solar wind. The software is designed around the principle that students can learn more by doing rather than by reading or listening. The programs provide a laboratory-like environment in which the student can control, observe, and measure complex behavior. The interactive graphics environment allows the student to visualize the results of his or her experimentation and to try different parameters as desired. The current version of the software runs on UNIX-based operating systems in an X-Windows environment. It has been used in a classroom setting at both UCLA and the University of California at San Diego.

1. INTRODUCTION

Laboratory courses traditionally have been an essential adjunct to classroom instruction in the physical sciences. Many principles are abstract until the student can experience the effects of those principles in a hands-on experiment. Moreover, many students learn better by doing rather than by reading or listening. The laboratory course strongly reinforces the book learning. In space physics, it is difficult to construct a traditional laboratory course because the size and expense of the apparatus necessary to undertake properly scaled experiments would be prohibitive. Fortunately, computers provide us with a mechanism to supplement the lecture course in a meaningful way. To support the classroom instruction in space physics at UCLA, we have developed a series of software modules that illustrate various concepts related to space plasmas. At this writing, we know of no similar available software for space physics. These modules provide an experimental experience with physical processes that cannot be studied in the laboratory. They allow the student to compare theoretical and model results with "observations" obtained using these modules and to visualize complex physical processes. They reinforce the classroom instruction and develop physical intuition.

Five modules are presently available. While designed as an adjunct to a lecture course in space physics, some of these modules would be useful in teaching more basic plasma physics concepts. The magnetospheres module takes the student through various magnetic field models of the Earth's magnetosphere and compares the dipole fields of each of the magnetized planets. The particle tracing module allows the student to follow the motion of ions and electrons of varying energy in magnetic

and electric fields of varying geometry. The cold plasma wave module allows the student to examine the behavior of electromagnetic waves in a cold plasma. The solar wind module illustrates the radial variation of the solar wind and interplanetary magnetic field, how the solar current sheet affects the structure of the interplanetary magnetic field, and how disturbances propagate through the interplanetary medium. Finally, the collisionless shock module allows the student to calculate how the solar wind plasma and magnetic field vary across a collisionless shock as the parameters of the shock and the solar wind vary. These modules will continue to be enhanced and new modules added during the foreseeable future. In particular, we expect soon to develop a current system module that illustrates the effect of currents flowing in different parts of the Earth's magnetosphere.

The development of software such as this is a complex and expensive undertaking. The programs described below were developed beginning in the summer of 1991 with the support of the National Science Foundation and NASA under the Space Grant Universities Program. In the following sections, we describe the programming environment and then each of the existing modules in turn.

2. PROGRAMMING ENVIRONMENT

The Space Physics Education Software was developed to run on Sun Sparc workstations using C, X-Windows, and Openlook. The software can be run "remotely" at our site at UCLA and displayed on any suitable local machine's display. The only requirement is that the local machine runs X-Windows. This includes Macintoshes, which run the Mac-X program. Local copies of the software can be made easily as well. Only two files need revision to create an executable on a new system. These two files are the Makefile and the X-Windows resource file. Please contact the lead author (ctrussell@igpp.ucla.edu) for information on using the code remotely or how to obtain a copy of the code.

Portability was kept in mind during development, so the program can run on different UNIX platforms and X-Windows environments. Currently, the software can be run on Sun workstations with either Sun's or GNU's C compiler for added

portability. Plans are to add Motif support to the program soon and to port the program to other vendor platforms as needed.

3. THE LEARNING MODULES

3.1 Magnetospheres

The magnetospheres module was designed to introduce students to some of the elementary properties of planetary magnetic fields. The first exercise introduces the *dipole magnetosphere* of each of the planets. The student can "fly through" the magnetic field using the mouse to take measurements at any chosen position. The next exercise allows the student to examine the properties of the "*mirror-dipole*" magnetosphere of Chapman and Ferraro and to see in a simple analytic model how the magnetosphere is compressed by the solar wind. The next level of complexity is the *spherical magnetosphere*, in which the equatorial field near the boundary is tripled because of the highly curved spherical magnetopause surface. This can be compared in the next exercise to *Tsyganenko's [1989a] vacuum magnetosphere* which is confined inside an elongated ellipse. With its realistically shaped dayside magnetopause, the field at the nose in this model is compressed by a factor of 2.44 over the undisturbed dipole field at this distance. The final exercise allows the student to probe *Tsyganenko's [1989b] empirical magnetosphere*, which is distorted relative to the vacuum magnetosphere because of the implicit presence of internal plasma. Figure 1 shows a copy of the window displaying this last option, with a train of mouse-selected position points shown on a "flight" through the magnetotail.

3.2 Particle Motion

The particle motion module was designed to demonstrate the behavior of single charged particles moving in magnetic fields of geophysical interest. On entering the module, the user is presented with a menu of field geometries, including *Uniform*, *Harris Sheet*, *Dipole*, *Mirror*, *Gradient*, and *Curvature* magnetic geometries and an *E×B* option that includes a uniform electric field together with a uniform magnetic field. The Harris current sheet is




Figure 1. Display from the magnetospheres module of the space physics program. The empirical model of Tsyganenko is shown. The x's are points at which the user has made measurements of the magnetic field using the cursor and the mouse. The results of the last measurement are given in the boxes on the left side of the display.

a simple analytic current sheet of finite thickness that approximates Earth's plasma sheet [Harris, 1962]. The user can then choose the desired particle mass (and charge), including H^+ , He^+ , He^{++} , O^+ , H^- , and "heavy" electrons. The particle can be started anywhere within the box showing the field geometry in three views by the manipulation of slider bars. Slider bars are similarly used to select the initial velocity vector components. The particle trajectory is then calculated by a standard finite difference algorithm for initial value problems. The tracing is terminated by use of a "pause" button. The trajectory can then be erased with another button and restarted as desired, or new trajectories can be superposed on the old using different selectable colors. Figure 2 shows a display for the dipole field option.

As the trajectory is calculated, the display in the upper right-hand corner shows the constantly updated particle energy, the accumulated time, the first and last pitch angles, and minimum and maximum positions and velocities. These allow the user to carry out quantitative "experiments," such as verifying the conservation of energy and adiabatic invariants, determining how mass and charge

affect drift motions, and so forth. In some options, parameters of the field model can be varied (e.g., the Harris Sheet thickness) so that the user can also obtain a feeling for how field strength, scale sizes of gradients and current sheets, and other modifications can affect the particle motion. When the particle motion has stopped, the user can measure positions on the screen with the mouse-driven cursor. These measurements appear at the bottom of the screen.

3.3 Solar Wind

The solar wind module was created to communicate concepts related to solar wind and interplanetary magnetic field behavior. On entering this module, the user chooses either the *Parker Spiral* or *Heliospheric Current Sheet* option. The *Parker Spiral* option allows the user to observe how radial motion of solar wind plasma can lead to the spiral interplanetary field geometry [Parker, 1958]. The user can use the slider bars to observe how the field geometry is affected by varying the solar wind velocity (or solar rotation rate) in either the equatorial plane or at a user-selected heliolatitude. The

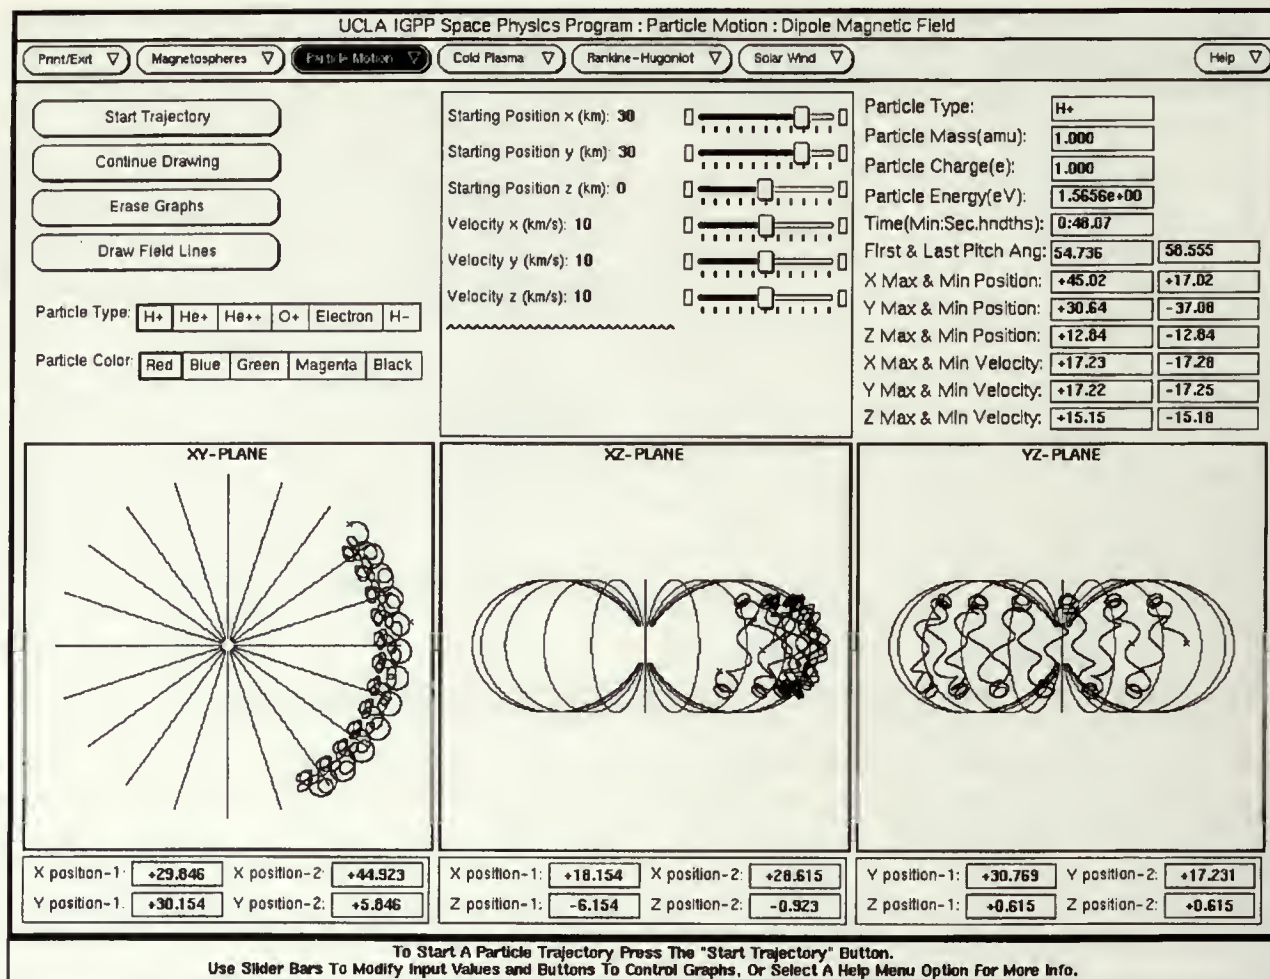


Figure 2. Display from the particle motion module of the space physics program. The particle-tracing exercise in a dipole field is shown. The boxes in the upper right update as the particle moves. The boxes on the bottom show positions measured by the user using the mouse and cursor.

garden hose angle and field strength and components also are displayed as a function of heliocentric distance. Planet locations are indicated on those displays to give the user a sense of the radial evolution of solar wind properties. Figure 3 illustrates the screen for the three-dimensional Parker spiral choice.

The *Heliospheric Current Sheet* option allows the user to "design" a heliospheric current sheet shape by combining magnetic axis tilt with the introduction of a quadrupolar contribution to the solar dipole magnetic field. The three-dimensional interplanetary current sheet shape is then computed and displayed at a user-selected perspective. Solar wind velocity can be varied by use of a slider bar. The displays together illustrate how the "ballerina skirt" model of the heliospheric current sheet is controlled by the shape of the neutral line at the solar "source surface." The associated

Stream Interaction option allows the user to impose a specific heliomagnetic latitude profile of solar wind velocity. It directs the program to compute an approximate model [Hakamada and Akasofu, 1982] of the distortion of the equatorial interplanetary magnetic field, given that velocity profile and the shape of a user-specified current sheet at the source surface. Associated model "time series" of solar wind properties at various heliocentric distances also can be displayed.

3.4 Cold Plasma Waves

The cold plasma waves module was designed to introduce students to electromagnetic waves in a cold plasma [Stix, 1962]. On entering the module, the user is presented with a menu of wave properties: *Index of Refraction*, *Dispersion Relation*,




Figure 3. Display from the Parker Spiral option of the solar wind module. Slider bars allow the user to change the velocity of the solar wind, the solar rotation rate, and the heliolatitude. The maximum radius slider bar allows the user to select the range covered by the bottom two panels.

Phase Velocity, Group Velocity, Parallel Group Velocity, Perpendicular Group Velocity, Ellipticity, and Wavelength. The user can then choose to display how a particular wave property varies as a function of either the frequency or the propagation angle (the angle between the wave vector and the background magnetic field).

Figure 4 shows an example of a “Phase Velocity” option screen. Slider bars are provided in the screen display that allow the user to select the background magnetic field strength and the plasma density. To view how the wave property varies as a function of these parameters, the user can select new values by moving the slider bars and pressing the “Draw Graph” button. Different colors are used to represent different wave modes (left-handed or right-handed). If the “NORMALIZED” button is chosen, the frequency is normalized by the proton

cyclotron frequency, the velocity is normalized by the Alfvén velocity, and the wave vector is normalized by the inverse of the ion inertial length (the velocity of light divided by the ion plasma frequency in radians per second).

To view how the wave property varies as a function of a variable, such as propagation angle, the user can select a frequency by either clicking the mouse button in the lower-right box or manually typing in the upper-left box, and then pressing the “Draw Graph” button. The lower-left box displays the results in a polar plot with the background magnetic field in the vertical direction. The user also can make a single-point measurement for any desired frequency, propagation angle, and plasma conditions. The wave properties displayed in the upper-right box correspond to the parameters in the upper-left box.

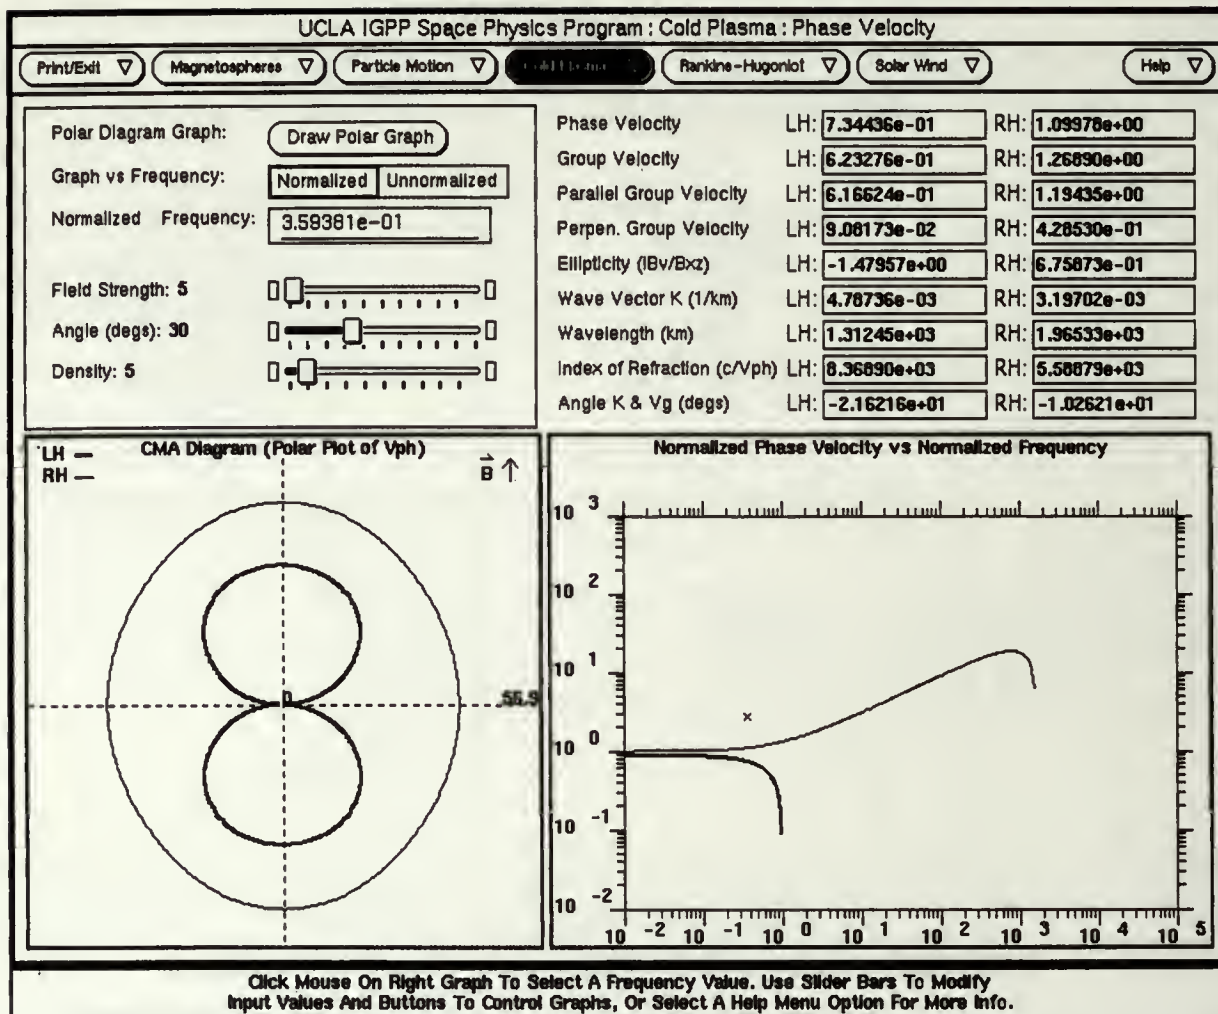


Figure 4. Display from the Phase Velocity option of the cold plasma module. The panel on the left shows a polar diagram of the phase velocity of left-handed and right-handed waves as a function of the angle of the wave normal to the magnetic field. The right-hand panel shows the phase velocity at fixed angle of propagation as specified by the slider bar versus frequency, here normalized by the proton gyrofrequency.

3.5 Collisionless Shocks: Rankine-Hugoniot Relations

The Rankine-Hugoniot module was designed to illustrate how the properties of a plasma, such as density, temperature, and magnetic field, change across collisionless shocks. The Rankine-Hugoniot conservation relations, which are incorporated into this module, allow for predictions of the properties of the downstream plasmas to be made based on knowledge of the strength of the shock and upstream conditions [Tidman and Krall, 1971]. In the *Graphs* section of this module, illustrated by Figure 5, the jump in number density, magnetic field strength, temperature, and plasma beta can be

calculated as a function of one of the controlling upstream parameters. One can vary the Mach number that measures the strength of the shock, the plasma beta that measures the ratio of the thermal to magnetic pressure, the angle between the upstream field and the shock normal, and the polytropic index in the ideal gas law. By selecting *None*, the user can find discrete values for the jumps in the plasma quantities for a given set of upstream parameters. The *Case Studies* section of this module allows the user to enter dimensional quantities for the plasma, such as velocity, number density, magnetic field strength, temperature, and so on, and obtain the downstream values for these quantities.




Figure 5. Display from the Graphs option of the Rankine-Hugoniot module. Here, the user has selected to display the variation of shock jump parameters versus Mach number.

4. OTHER SOFTWARE USED IN UCLA SPACE PHYSICS COURSES

Interactive, menu-driven graphics software is used both as a tool for graduate students in their dissertation research and also in computer laboratory exercises to introduce students to the physical processes and phenomena in space plasma physics. These software tools allow immediate access and display of the data and facilitate the application of standard analyses to the data, such as minimum variance analysis; Fourier analysis, and filtering [Russell, 1983].

Modern commercial mathematical packages are now available that facilitate the manipulation of mathematical expressions, the integration of functions, the solution of differential equations, and the display of the results of these manipulations. At UCLA, we have used *Maple* (Waterloo Maple Software, info @ maplesoft.on.ca) in computer laboratory exercises and find that students often

extend their use of this software beyond the classroom setting.

5. CONCLUDING REMARKS

From our experience at UCLA, interactive, menu-driven graphics software is a good way to introduce students to the physical processes occurring in space plasmas. We have developed modules for magnetospheres, particle motion, cold plasma, solar wind, and collisionless shocks. These modules need extension, and we need more modules to provide a more complete curriculum. They are now being used in both upper division and graduate classes at UCLA and have recently been introduced at the University of California at San Diego. The modules have been best received in computer laboratory situations where the instructor is available to answer questions. Remote dial-in usage has been less successful principally because of interfacing problems with X-Windows emulators. Gradu-

ate students prefer remote dial-in capability because it allows them the freedom to arrange their schedules, but such freedom comes at the price of decreased interaction with the instructor.

We welcome other users. We also welcome new ideas for modules, and, especially, we welcome assistance in developing modules. Nevertheless, despite our success to date, some problems remain. First, developing software is expensive. Second, because graduate students and outside users want to run software on a variety of platforms, portability is critical. Fortunately, current developments in computer software and operating systems may assist in mitigating, if not completely solving, these two problems in the coming years.

Acknowledgments. We wish to acknowledge the help of Marilyn Van Swol in developing an early version of this program. We are also grateful to H. Herbert who has helped install and provide access to this software, to S.M. Petrinec and M.H. Farris who developed algorithms for two of the modules, and to all the students who have provided feedback on the use and functionality of the programs. This work was supported by NASA under research grant NGT-40005 and by the National Science Foundation under research grant USE 91-55988.

REFERENCES

- Hakamada, K. and S. I. Akasofu, Simulation of Three-Dimensional Solar Wind Disturbances and Resulting Geomagnetic Storms, *Space Sci. Rev.*, 31, 3-70, 1982.
- Harris, E.G., On a Plasma Sheet Separating Regions of Oppositely Directed Magnetic Field, *Nuovo Cim.*, 23, 115, 1962.
- Parker, E.N., Dynamics of Interplanetary Gas and Magnetic Fields, *Ap. J.*, 128, 664-676, 1958.
- Russell, C.T., Interactive Analysis of Magnetic Field Data, *Adv. Space Res.*, 2(7), 173-176, 1983.
- Stix, T.H., *Theory of Plasma Waves*, McGraw-Hill, New York, 1962.
- Tidman, D.A. and N.A. Krall, *Shock Waves in Collisionless Plasmas*, Interscience, New York, 1971.
- Tsyganenko, N.A., A Solution of the Chapman-Ferraro Problem for an Ellipsoidal Magnetopause, *Planet Space Sci.*, 37, 1037, 1989a.
- Tsyganenko, N.A., A Magnetospheric Magnetic Field Model With a Warped Tail Current Sheet, *Planet Space Sci.*, 37, 5, 1989b.

Data Base Management System and Display Software for the National Geophysical Data Center Geomagnetic CD-ROMs

N.E. PAPITASHVILI¹

*National Space Science Data Center, NASA Goddard
Space Flight Center, Greenbelt, MD*

V.O. PAPITASHVILI²

*Space Physics Research Laboratory, University of
Michigan, Ann Arbor, MI*

J.H. ALLEN and L.D. MORRIS

*National Oceanic and Atmospheric Administration,
National Geophysical Data Center, Boulder, CO*

The National Geophysical Data Center has the largest collection of geomagnetic data from the worldwide network of magnetic observatories. The data base management system and retrieval/display software have been developed for the archived geomagnetic data (annual means, monthly, daily, hourly, and 1-minute values) and placed on the center's CD-ROMs to provide users with "user-oriented" and "user-friendly" support. This system is described in this paper with a brief outline of provided options.

¹ On leave from WDC-B2/Geophysical Center, Moscow, Russia.

² On leave from IZMIRAN, Troitsk, Moscow Region, Russia.

The National Geophysical Data Center (NGDC) has collected geomagnetic data from the worldwide network of magnetic observatories for most of the last four decades. NGDC has accumulated the world's largest collection of magnetograms on microfilm and microfiche, publications, hourly mean value tables, and digital geomagnetic data on various media, including magnetic tapes, optical erasable discs, and CD-ROMs. The NGDC CD-ROMs with solar-terrestrial physics data were issued in 1987 and 1990. The data base management system and retrieve-and-display software have been developed for these CD-ROMs, including an access to all kinds of the archived geomagnetic data (annual means, monthly, daily, hourly, and 1-minute values) to provide users with "user-oriented" and "user-friendly" support. The current version of software is oriented to support the hourly and 1-minute data bases, and it provides daily, monthly, and annual averages from the hourly mean values (Figure 1).

The "hourly" geomagnetic software does the following:

- Displays the geomagnetic observatory catalog [Abston et al., 1985]—observatory name, three-letter observatory code, two-letter country code, and geographic coordinates (Figure 2)
- Displays the monthly data availability information table—that is shows the years and months of records of a given observatory that are available in the hourly means digital data base, such as CD-ROM (Figure 3)
- Creates a subset of selected data in a flat ASCII file (Figure 4) and provides an option to sort the data from the standard

- “1 month/one component” sequence to the “1 day/all components” sequence, whichever is more convenient for scientific analysis
- Displays the hourly data file day-by-day or provides multi-day graphics (up to 31 days)—the latter option provides a monthly “stack-plot” magnetogram of all components for a given month (Figure 5), and 31, 30, 29, or 28 days will be shown for the corresponding months
- Allows changes to the graphic and background colors for the convenience of the output to a color printer, shows an average value of each component for the selected interval, allows changes to the scale factor of the components, and views the file forwards and from the beginning
- Retrieves the hourly mean values from the original file to compute the daily and monthly mean values (in the separate flat ASCII files) and displays daily values for the entire year to show the annual variation of the geomagnetic field (Figure 6)—this option helps carry out quality control of geomagnetic data and conduct the scientific analysis of secular variation of the geomagnetic field

The “1-minute” *geomagnetic software* retrieves and creates a subset of selected data in a flat ASCII file and shows the data in a 24-hour group (Figure 7). The graphic mode provides the same visualization service as described above for the hourly mean values.

The *retrieval/display software* has been developed as a set of independent executable modules, which can be easily replaced by new ones. Therefore, it is easy to add new options for the data visualization service. The Geomagnetic Data Base Management System consists of the following independent modules:

- Main data base management system program (run first)
- Hourly value retrieval program
- Module to sort hourly value data in the sequence “all components for 1-day”
- “Hourly averaged” daily variation graphic program (displays data from 1 through 31 days on the screen)

- “Daily averaged” annual variation graphic program (daily mean values are read from the hourly value data file)
- One-minute value retrieval program
- “One-minute averaged” daily variation graphic program
- Documentation file

All programs are written in MS-DOS C language (Version 6.00A) and may be run on the widely used 286, 386, or 486 personal computers. Retrieval modules create intermediate files INS and INSM in the current directory, which contains the names of the created output files (subsets of the hourly value and 1-minute value data, respectively). The graphic visualization modules use these files to display the hourly value and 1-minute value daily or annual variations. These graphic modules can be used separately from the data base if the names of displayed data files are placed in the INS or INSM files by the standard text editor. File LCOLOR is created from the first run of any graphic program. The file contains the number “1” to plot the color graphics over the black background of the screen. By changing this number to “2,” the color graphics will be plotted over the white background; “0” displays the white graphics on the black background. This option is a menu choice, but the numbers also can be placed in the LCOLOR file by the standard text editor.

The hourly values must be written in the IAGA format—120 ASCII characters per logical record, or 62 binary bytes of the NGDC-05 CD-ROM format. The “hourly” software is used for the current management of the NGDC digital hourly value geomagnetic data base on CD-ROMs and optical-erasable discs. The “1-minute” data must be written in the WDC-A format—400 ASCII characters per logical record.

The “1-minute” software has been selected as a primary software set for future NGDC geomagnetic CD-ROMs, which will contain 1-minute geomagnetic data collected in the framework of STEP Project 6.4, “Global Geomagnetic Database for STEP.” The first NGDC/STEP CD-ROMs could contain 1-minute geomagnetic data for 1990-1991 from up to 50 to 60 magnetic observatories worldwide [Papitashvili et al., 1993]. Eventually, a set of CD-ROMs will cover the entire STEP period, 1990 to 1997. This effort is consistent with the

approach to bring the "Personal Data Center" on CD-ROMs (i.e., the World Data Center products) to scientists in any corner of the world. The NGDC CD-ROMs and accompanying software are available from the National Geophysical Data Center (Boulder, CO) at a minimum cost to cover only shipping and handling.

REFERENCES

Abston, C.C., N.E. Papitashvili, and V.O. Papitashvili, *Combined International Catalog of Geomagnetic Data*, Report UAG92, 291 pp., NOAA/National Geophysical Data Center, Boulder, CO, 1985.

Papitashvili, V., M. Teague, E. Friis-Christensen, and J. Allen, STEP Project 6.4 Plans to Establish One-Minute Geomagnetic Database, Part 2, *STEP International*, 3, 1, 6-9, 1993.

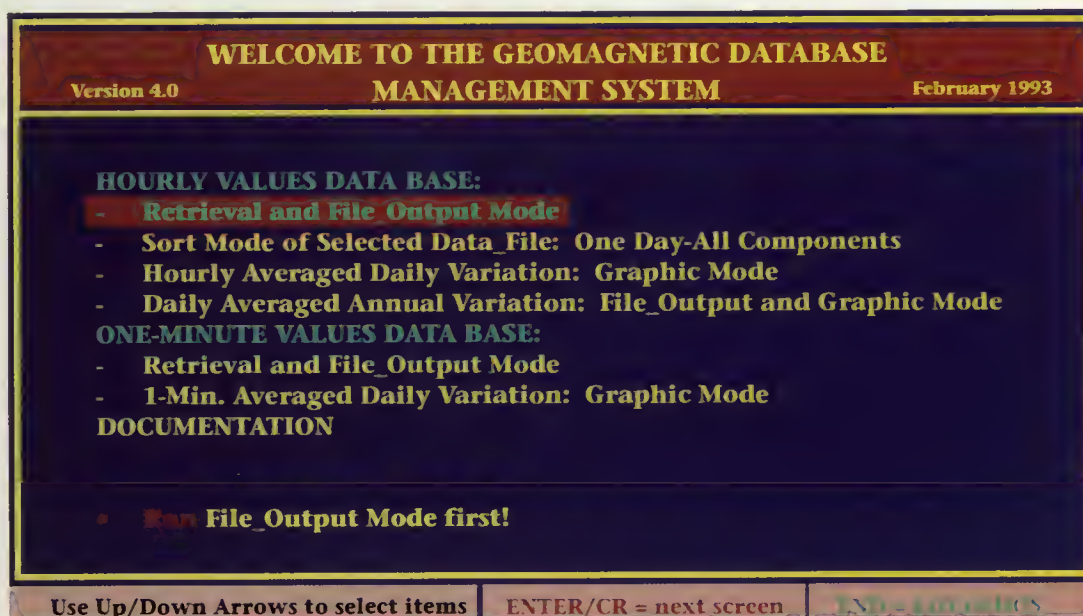


Figure 1. Geomagnetic Data Base Management System main menu.

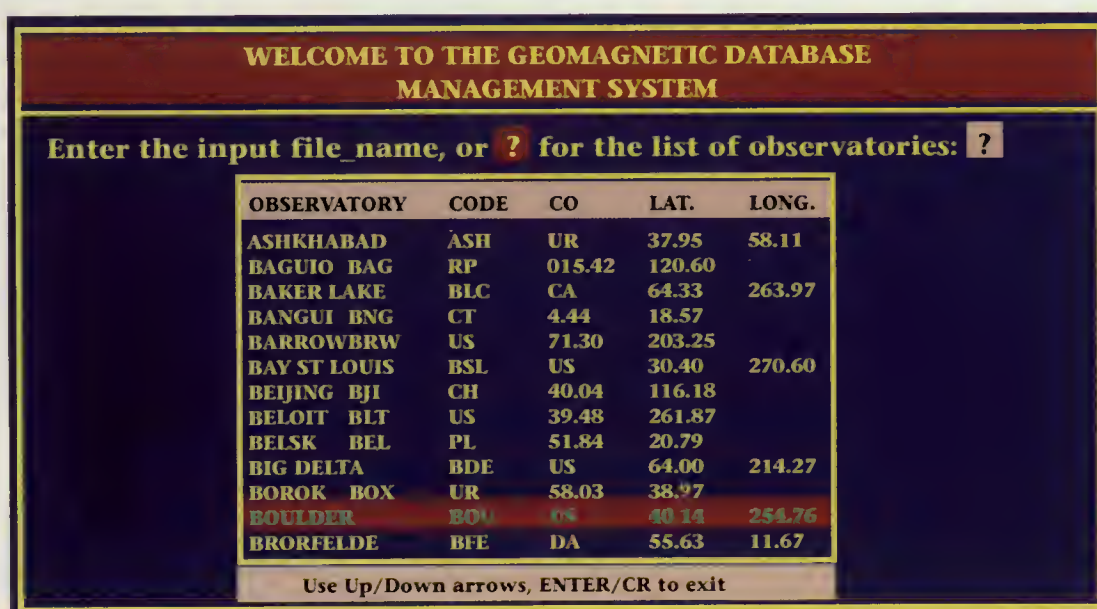


Figure 2. Listing of the observatory names and three-letter codes, two-letter country codes, and geographic coordinates.

**WELCOME TO THE HOURLY VALUES DATABASE
RETRIEVAL SYSTEM**

Enter the input file_name, or ? for the list of observatories: **F:BOUHVPF**
 Enter the output hourly values data file_name: **BOU 1989**
 Enter **Y/y** for CR/LF at the end of output logical record, or **N/n** : **n**
 Observatory **BOU** data_file contains the following time_interval:

BEGINING	END
Year=67	Year=92
Month=01	Month=12
Day=01	Day=31
Components: DHZ	

Observatory BOU		
1980	JFMAMJ	JASOND
1981	JFMAMJ	JASOND
1982	JFMAMJ	JASOND
1983	JFMAMJ	JASOND
1984	JFMAMJ	JASOND
1985	JFMAMJ	JASOND
1986	JFMAMJ	JASOND
1987	JFMAMJ	JASOND
1988	JFMAMJ	JASOND
1989	JFMAMJ	JASOND

Press ENTER to exit

Figure 3. Monthly data availability information table.

**WELCOME TO THE HOURLY VALUES DATABASE
RETRIEVAL SYSTEM**

Enter the input file_name, or ? for the list of observatories: **F:BOUHVPF**
 Enter the output hourly values data file_name: **BOU 1989**
 Enter **Y/y** for CR/LF at the end of output logical record, or **N/n** : **n**
 Observatory **BOU** data_file contains the following time_interval:

BEGINING	END
Year=67	Year=92
Month=01	Month=12
Day=01	Day=31
Components: DHZ	

Hit SPACE to read all components consequently or enter one only :

Enter time_interval to retrieve data:

BEGINING	END
Year=88	Year=88
Month=01	Month=12
Day=01	Day=31

Figure 4. Time interval selection menu.




Figure 5. Hourly value "stack-plot" magnetogram of the Boulder magnetic observatory for October 1989.




Figure 6. Daily averaged annual geomagnetic field variation—Boulder magnetic observatory, 1989.




Figure 7. One-minute geomagnetic data display—Boulder magnetic observatory, October 21, 1989.

Visual Interface for Space and Terrestrial Analysis

EDMUND G. DOMBROWSKI

SFA, Inc., Landover, MD

JASON R. WILLIAMS and ARTHUR A. GEORGE

Hughes STX, Lanham, MD

**HARRY M. HECKATHORN and
WILLIAM A. SNYDER**

Naval Research Laboratory, Washington, DC

The management of large geophysical and celestial data bases is now, more than ever, the most critical path to timely data analysis. With today's large volume data sets from multiple satellite missions, analysts face the task of defining useful data bases from which data and metadata (information about data) can be extracted readily in a meaningful way. Visualization, following an object-oriented design, is a fundamental method of organizing and handling data. Humans, by nature, easily accept pictorial representations of data. Therefore graphically oriented user interfaces are appealing, as long as they remain simple to produce and use. The Visual Interface for Space and Terrestrial Analysis (VISTA) system, currently under development at the Naval Research Laboratory's Backgrounds Data Center (BDC), has been designed with these goals in mind. Its graphical user interface (GUI) allows the user to perform queries, visualization, and analysis of atmospheric and celestial backgrounds data.

1. INTRODUCTION

1.1 The Data Center Challenge

Geophysical and astronomical data centers throughout the world are currently facing the challenge of how to properly manage large-volume data from multiple satellite missions in an efficient manner. With the advent of remote sensing instruments producing gigabytes of data a day, both computer scientists and analysts are left with the realization that data collection technology is outpacing data archival and access technology. Scientists, programmers, and data base analysts are therefore presented with the problem of defining more useful scientific data bases from which one can readily extract data and metadata (information about data) in a meaningful way.

At the Backgrounds Data Center (BDC), one of the main charters is to catalog, access, and view mission information pertinent to each experiment data base. BDC is responsible for handling DoD and non-DoD remote sensing terrestrial and celestial backgrounds data. A majority of the software development time at the data center involves the design and implementation of on-line catalog systems. The initial catalogs have been built using a relational data base management system (RDBMS) called INGRES®. Relational tables containing critical experiment metadata are created and manipulated through the use of standard Structured Query Language (SQL) commands using a forms-based user interface (FBUI).

An example of an online catalog is the BDC Summary Catalog, which enables a user to build temporal, spatial, and spectral (TSS) queries, run them against all available INGRES tables, and view the resultant items in a list box. A user works

through a series of text windows, using the keyboard and specific function keys to build the query requirements. The idea at the summary catalog level is to help a user properly evaluate the data sets that match certain criteria and to even perform a basic cross correlation of results. The Summary Catalog, however, only addresses the metadata of these experiment data bases, not the actual data. It contains mission-summary-level information, such as platform instrumentation specifications, spacecraft positional information, and observation events for each BDC mission. Data are not directly accessible. However, information relative to the actual data is provided, including a reference to the physical archive so a user can place an order for the actual data items.

For more specific information about experiment data, BDC created detailed index catalogs for certain missions. These catalogs, called Program Catalogs, enable a user to access lower level metadata (data item specific) using a series of FBUI text windows in a fashion similar to the Summary Catalog. However, just like the Summary Catalog, the Program Catalogs only reveal detailed metadata results from data, not the actual data. These catalogs are sufficient for certain purposes, such as mission correlation and analysis of indexed results versus time. However, if a user wants to view metadata information and resultant parameters (e.g., graphs, line plots, and graphic windows) or work with the actual data, other tools must be used.

1.2 The VISTA System

BDC began developing a tool to provide visualization of satellite viewing scenarios. The tool grew into a more sophisticated system that not only provided a graphical view, but used the graphics to actually build, run, and view query results from the catalogs. The Visual Interface for Space and Terrestrial Analysis (VISTA) took a step beyond the FBUI methods of catalog access by including visualization and direct data access. VISTA provides a visual query tool for access to the BDC data bases. This, in itself, is a distinct advantage over the more traditional FBUI catalogs, but the functionality of VISTA did not end there. VISTA also added the ability to directly access data as selected from metadata lists within the same interface. Hence,

VISTA achieved the goal of being a fully integrated catalog system—that is, able to access and visualize both metadata and data.

Visualization, using an object-oriented design, is a fundamental way of organizing and handling metadata and data. Humans, by nature, assimilate data more rapidly by visual means than any other. VISTA has been designed to meet the challenge of bringing graphical catalog access and data visualization to users. This design has resulted in a graphical system pertinent to geophysical and celestial backgrounds data in which direct visualization and information about the data are put in the context of how they were obtained. In addition, the VISTA system has also met the challenge of how to efficiently manage the metadata and data files themselves. The VISTA design, interface and method of data base management is discussed in this paper. For additional background material on VISTA and BDC, refer to Dombrowski et al. (1992a, 1992b, 1993, 1994) and Snyder et al. (1991, 1993).

2. VISTA DESIGN

2.1 Software Design

The VISTA prototype system was developed under the X11 Window System (OSF/Motif) using the C programming language on a UNIX platform. The majority of the VISTA windows are strictly X-Windows. However, because the prototype development was initiated on Silicon Graphic, Inc. (SGI), workstations, the current interface also makes use of GL windows for some of the two- and three-dimensional interactive windows.

The system architecture consists of several major software layers. The top layer, User Interface, makes full use of the X11 Window System using the Motif window manager and C. SGI's GL is also used, as mentioned above, but these windows are encapsulated within X-Windows frames. The User Interface is described in greater detail in section 3. The second layer is Session Management, which handles all sessions parameters such as user identification, user area access, metadata file location, and pointers to system files. The third layer, Access, handles two main object types: User Management Objects and Query Objects. The

Management Objects consist of files created as a result of a user session. The Query Objects are parameters that are built and saved as part of the user query. The bottom layer, Data Management, is split into two separate systems: an RDBMS and a data structure flat file management system. Both systems use the same top three levels. However, they are both uniquely tied into the VISTA interface system.

2.2 Data Management Systems

The VISTA system has been developed with two optional data base management systems: an SQL/INGRES-based RDBMS and a data structure flat file management system. The latter is the portable version that has been the focus of the prototype system. The former is one that works more specifically for BDC RDBMS catalog tables, but it also can be used to support similar structures elsewhere. In fact, it is the Catalog Link section of VISTA (see section 3.2) that should expand the use of this version because a multitude of useful data bases now exist under an RDBMS. Therefore, at other data analysis centers, VISTA could be extended to support various data base management systems and the information they hold.

The portable version of VISTA currently works with data structures that make use of metadata flat files and directories containing the actual data online. Part of the design of the system allows the use of routines inherent in the INGRES-dependent version in retrieving and building the metadata flat files for the portable version. Therefore, additional mission information fed into the INGRES catalog systems at BDC or elsewhere can be queried and used to populate VISTA flat files for the portable version.

3. VISTA PROTOTYPE INTERFACE

3.1 The Main Window

The VISTA Main Window (see Figure 1) is the focal point of the system. The window displays four main user access areas: Catalog Link, Build Query (and Run Query), Edit/View, and Analyze. Each area is color coded to help the user move through the interface more rapidly. Multiple windows opened in an area are easily recognized by the color-coded bars at the top of the window. These four main VISTA areas are discussed in separate sub-headings in a fashion parallel to how a user would activate the interface.



Figure 1. VISTA Main Window.

A fifth area of the Main Window associated with the four main areas is the VISTA WorkSpace. This display window is the area where users can save particular objects (see Figure 1) as produced from one of the main areas. For instance, a user in Build Query may want to save the parameters created in a TSS query. Within Build Query, there are options to allow the user to select this save mechanism. A specific color-coded object corresponding to the saved query is placed in the VISTA WorkSpace area. This object contains all saved query parameters and is tagged with a specific filename for retrieval and use within the interface. Each object has an associated pull-down menu that allows the user to activate, rename, copy, or delete the object at any time. The VISTA WorkSpace area is similar to PC and workstation environments that rely on visual icons for files and executables. By using the WorkSpace, a user can enter and exit the VISTA system at any time while preserving the information created within a particular area of the Main Window.

3.2 *Linking to the Science Catalog*

The first button on the left of the Main Window is Catalog Link. In the prototype, this button is inactive, hence grayed-out. In future versions, this button will provide access to other catalog interfaces and be a mechanism for the retrieval of pertinent metadata from sources other than the VISTA internal data base. For instance, by pressing the Catalog Link button, users will attain direct entry into the BDC Summary and Program Catalog interfaces and the BDC Science Catalog Information Exchange System (SCIES), and they have direct access to Internet, FTP services, and hypertext tools such as Mosaic®, which was created by the National Center for Supercomputing Applications (NCSA).

3.3 *Building Queries*

The second button from the left, Build Query, provides direct access to a series of windows allowing a user to specify query parameters to run against various mission sources (data bases). When a user activates the Build Query button, the Build Query control panel appears (see Figure 2). This

panel consists of three main areas: Source List, Query Parameters, and Run Query. The Source List area consists of a box presenting all available data sets from which to query against. The list presents the names of the missions supported at BDC. From the list, a user can select either an entire mission or sub-select a particular instrument or observation, and so on, depending on the most logical characterization of the mission parameters. A user can also select multiple missions or portions of mission data for cross-correlation purposes. This is the real power of multiple or sub-sampling source selection list boxes, allowing a user to perform various cross-correlational studies.

The Query Parameters section of this control panel is where a user builds the detailed TSS query specifications using both list boxes and activation buttons. Queries are built by either directly typing them in or by using graphical windows to delineate specific query parameters. For each TSS specification, there is an associated window activation button. For example, for spatial specifications, a user activates either Geographical (Earth observing systems) or Celestial (space observing systems). These graphical windows allow the user to specify multiple spatial queries. The Geographical Query window (as shown in Figure 3) allows a user to specify Earth-based spatial ranges. The Geographical window provides a map view of the Earth and allows the user to move around, using slider bars, in latitude and longitude.

Various control buttons in the Geographical Query window enhance the area selection process. For instance, the user has the ability to zoom, change the map selection, or create a query area. When a user wants to create one or more areas in latitude and longitude, the Create Area button is selected. Input parameters from the geographical map view are saved in the GeoArea located below the map view (see Figure 3). If a user wants to specify an altitude range for the data query, the altitude bar to the right of the geographical map view is repositioned. This bar mechanism allows the user to set a low and high range in the atmosphere and to associate it with the GeoArea in latitude and longitude. By using this graphical tool, a user builds direct spatial queries for Earth-based observations from remote sensing platforms. When finished building queries, a user presses the "OK"

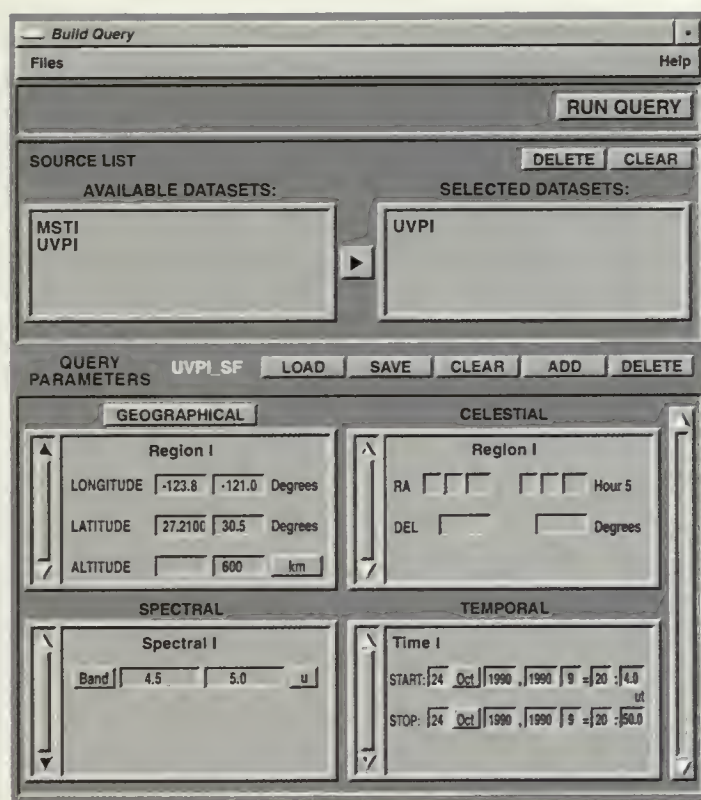


Figure 2. Build Query control panel.

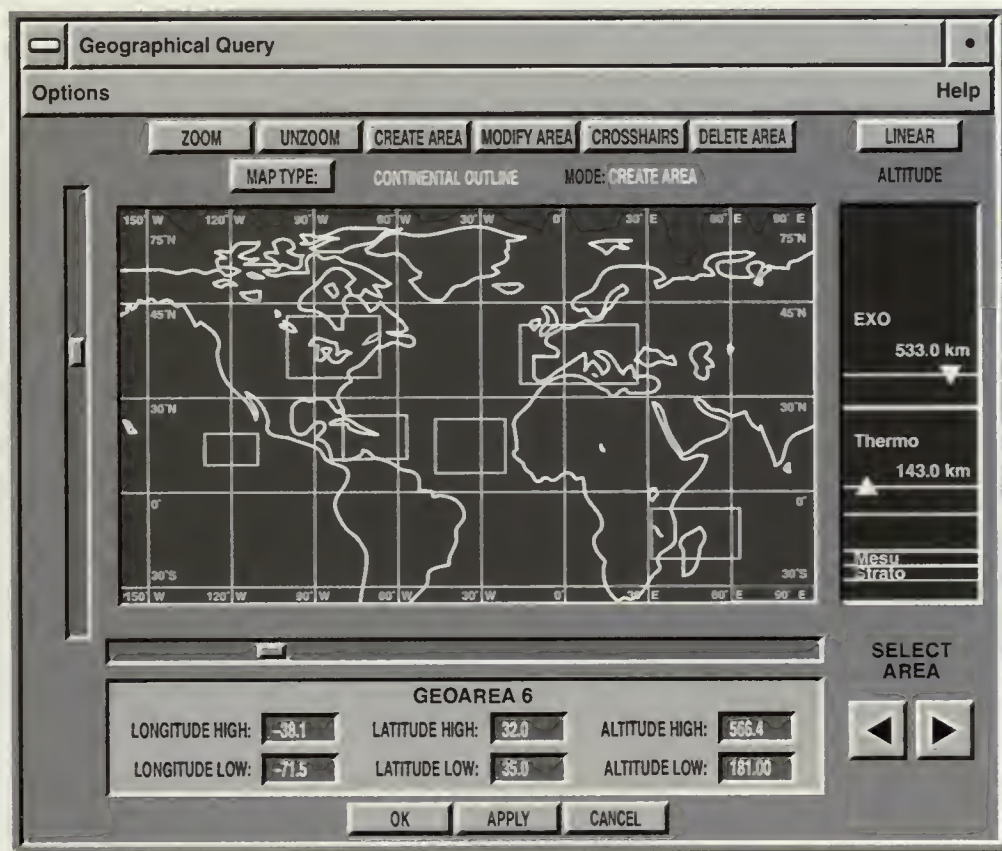


Figure 3. Geographical Query window.

button at the bottom of the Geographical window to update the Build Query control panel. Values can be changed in either the Geographical section of the Build Query control panel or the graphical windows.

The Query Parameters section of the control panel also allows a user to save the query parameters as new Workspace objects or to clear the control panel and bring in pre-saved Build Query objects. A user can create various combinations of query parameters and sources to run multiple queries against the catalog data bases. Allowing users to create multiple queries demonstrates the power of a graphical-based user interface, which makes data base query building faster and more efficient.

3.4 Running Queries

The Build Query control panel contains the windows for building queries and the Run Query button, which activates the query process. The VISTA prototype was initially built relying on an RDBMS query mechanism, making it more practical for BDC use because there is an INGRES catalog dependency. When this version is run, SQL commands are built and sent to query against the INGRES tables containing the appropriate VISTA metadata. At BDC, a user running the INGRES version of VISTA on a workstation would use the interface to build and run the query, and when he or she activates the run query, the session management layer of VISTA takes the SQL commands and sends them over the network to back-end main frame machines that host the INGRES tables. The information retrieved from the tables is then passed back to the individual workstation and VISTA, which stores the incoming metadata values into appropriate dynamic variables.

In addition to the INGRES-dependent version, VISTA also offers a completely portable non-INGRES version for the data base management system. With this system, users are only dependent on their workstation platform to run VISTA. In this case, when a user hits the Run Query button, a completely internal file management system is used to provide query results to the system. Currently this is the more mature system under development because it is this version that will be distributed for review and use. This version uses data structures to directly search indexed metadata flat files as opposed to relational tables.

3.5 Editing and Viewing Query Results

When a user presses the Run Query button in the Build Query control panel, the Edit/View area is automatically activated. The Edit/View control panel (see Figure 4) is the area that gives the user the first view of the metadata results. This panel is divided into two main sections: Data List (query results list box) and View List (selected data items for viewing). The Data list provides the user with a column-formatted table that displays the data items matched by the query and their associated key metadata values. A typical line in this list would display the associated source, the item catalog identification number, and key physical parameters such as positional and temporal information. Users can select individual items, groups of items, or all of the Data List return values to be put into the View List. Mechanisms to save and print the resulting lists are a planned VISTA enhancement.

The View List is specifically created for direct visualization of metadata items. The View List works with a series of graphical visualization windows specifically created for geographical- and celestial-based data. At the bottom of the Edit/View control panel are five window buttons: Map View, Vista View, Celestial, Line-of-Sight, and QuickLook. The Map View window is similar to the Geographical Query window. It provides a map view of the Earth for Earth-based observations, along with multiple option buttons for viewing the selected metadata items. This window gives the user a two-dimensional projection of the observing platform over the Earth, a projection of the FOV corner points, the line-of-sight vector, and the area covered in the atmosphere or on the Earth's surface as a geometric projection.

The Vista View window (see Figure 5) provides a similar view, but as a three-dimensional interactive window that allows the user to move around the Earth as if above the Earth in a "meta-observer" position. Movements include panning, zooming, changing latitude, longitude, and altitude, or changing the look point from the satellite's position. A position of the satellite is plotted, as well as similar information found in the Map View window. The user also has the option to change coordinate system plots or turn Earth lighting models on and off. The Vista View clearly offers some of the more powerful

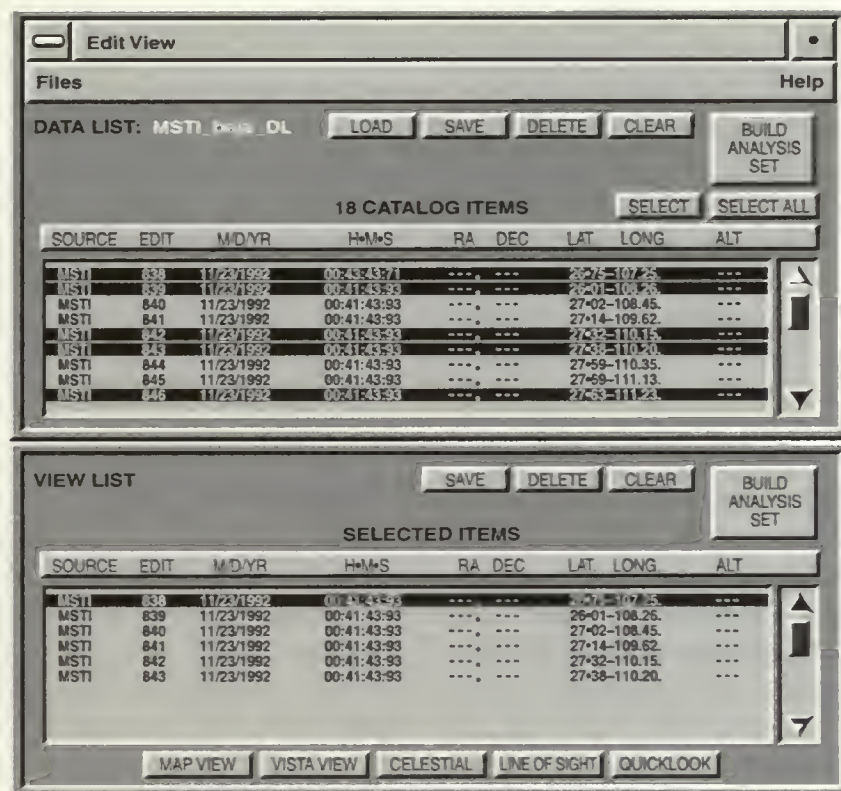


Figure 4. Edit/View control panel.




Figure 5. Vista View window.

tools for data set evaluation. A third window, directly tied to the previous two, is the Line-of-Sight window. This graphic is similar to the Vista View in that it gives you a three-dimensional view; however, this window was designed to give the user the sensor eye perspective of an observational event. The window reflects the actual position of the instrument above the Earth's surface and shows the view from the FOV dimensions.

The QuickLook and Celestial viewers are independent from the other three buttons (not available for prototype release). The QuickLook window will provide a representative item from each data set, such as an image, or spectra, or even line data, depending on the type of instrument used. The Celestial window will display a celestial background map using a celestial catalog of choice, and it will overlay the areas observed. Celestial object information will be available, such as location, magnitude, and object type for stars and extended objects.

All of the Edit/View windows have been designed to help users make decisions about certain data results by viewing both the metadata lists and the visualization at the same time. At any time, the user may save the Edit/View list as an object to reside in the Workspace area. A user may also save portions from these lists by using the Build Analysis Sets button to build analysis sets that contain pointers to the actual data.

3.6 Analyzing Data

After refining data lists in the Edit/View section of the interface, a user may want to make the next natural progression—that is, work with the actual data. The analysis sets that are built from the metadata in the Edit/View section are the vehicle for image processing and analysis within the system. Each analysis set, based on the list of selected data items, contains the necessary information for a user to pull up online data and to do this within an analysis tool without leaving the VISTA system.

When a user presses the Analyze button in the Main Window, the Analysis control panel (see Figure 6) is activated. This window lists the data items that were chosen for the particular analysis set and provides optional selections for the user, such as image processing tools, analysis packages, and various data modeling tools. The Analyze section of VISTA has been designed to be modular so each user can implement a set of favorite tools or incorporate user-developed tools into the VISTA framework. Standard proprietary data processing tools such as IDL®, IRAF®, SAO Image®, PV-Wave®, AVS®, Iris Explorer®, or Data Explorer® can have direct links built into the VISTA Analysis section. At BDC, direct links for several of these packages will be implemented. Sessions can be started and data loaded from within the Analysis control panel,

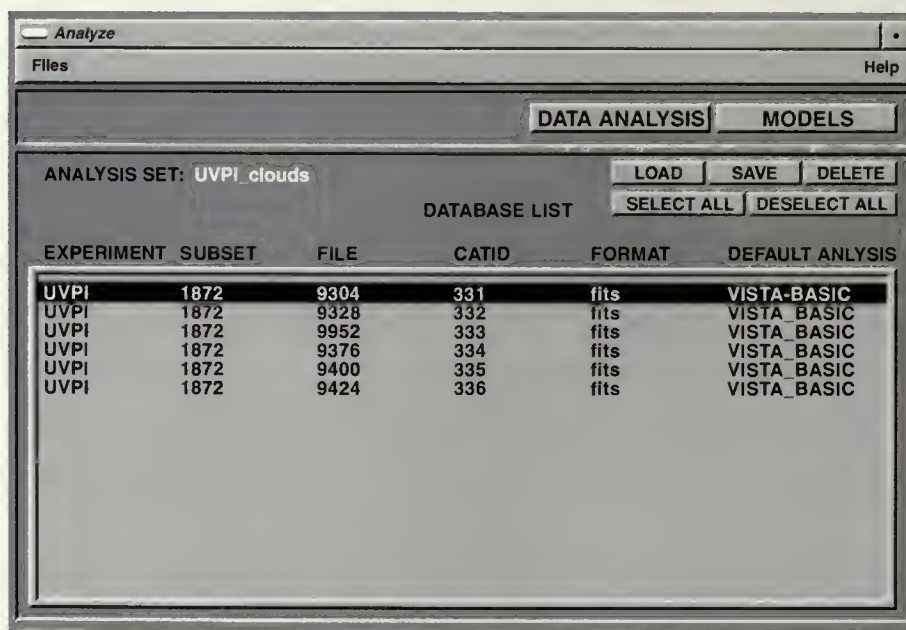


Figure 6. Analysis control panel.

making some of the often cumbersome initialization steps and data loading processes transparent to the user. Even better, non-proprietary tools such as Khoros (©Khoros Research, Inc., 1994) or NCSA routines can be packaged with VISTA and distributed to users. In fact, the VISTA prototype uses Khoros as its main image processing and analysis package. Khoros was initially developed at the University of New Mexico as a DoD project and since has become an open software package that users can get via Internet. Because BDC is a member of the Khoros consortium, Khoros programs can be developed by the VISTA team and distributed along with the package to users.

The VISTA prototype works with defined data processing pipelines that involve the Khoros tool, Cantata, which is a visual programming environment tool. This system presents a user with a workspace from which program pipelines can be built using available Cantata routines or user-developed routines hosted by Cantata. Cantata displays various image and data processing routines as glyphs (see Figure 7) that can be connected together to form pipelines. When run, each glyph is activated and passes on I/O to the next stage in the pipeline. VISTA uses a basic pipeline for data

display as its default Cantata pipeline. That is, when a user brings in a particular analysis set through the Analyze area and activates Cantata with the default pipeline, Cantata is automatically initiated and the pipeline and data selection loaded. At this point, the user is within the Cantata system without leaving VISTA. To run the pipeline, a user only has to click the run button on the left panel of Cantata. Figure 7 shows the Cantata interface, as brought up using VISTA, and shows the output of image display data for a particular analysis set.

VISTA also has been designed to enable users to make use of existing workstation tools (i.e., Image Tools or XV) or user-defined tools to perform screen captures (RGB, TIFF, Postscript, etc.) of any combination of VISTA and analysis plot windows. These screen captures can then be used to create hard-copy output or animation sequences for videos or workstation movie routines.

4. FUTURE ADAPTATIONS AND ENHANCEMENTS

The VISTA prototype was released as a beta version in October 1993. This elicited positive responses from the beta reviewers and provided the

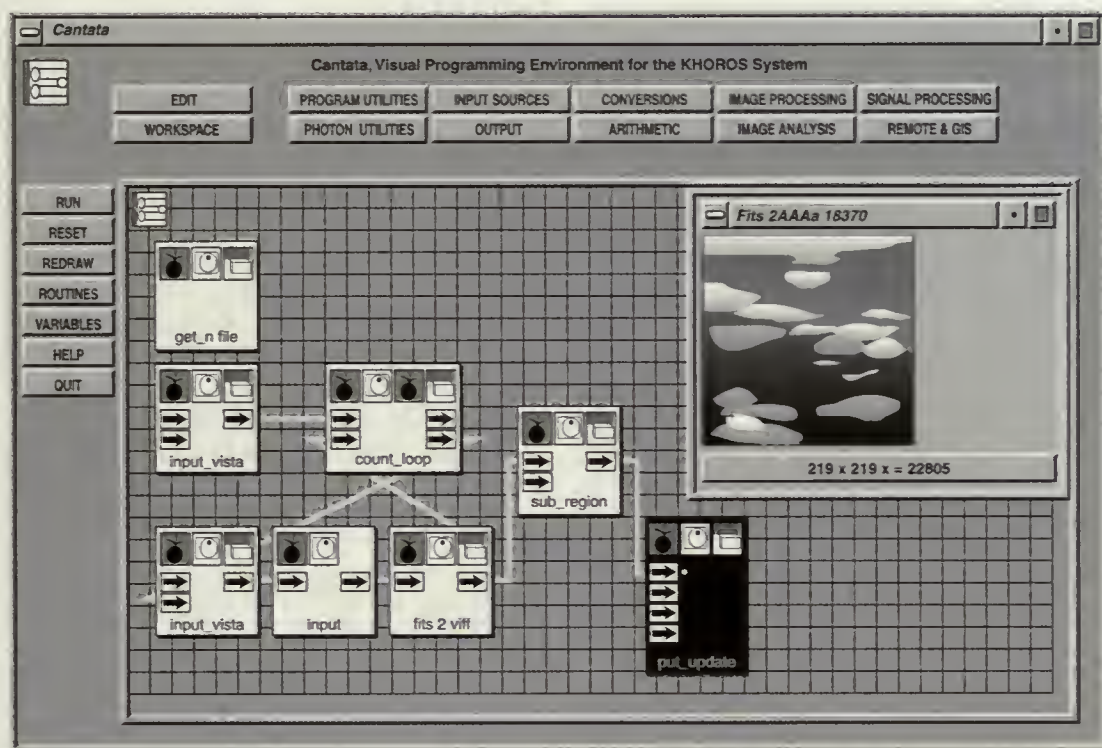


Figure 7. Cantata.

team with a basis for improvements for prototype updates. The next prototype version is targeted for release at the end of January 1994. While work continues on the prototype system, the VISTA team is also undertaking a formal object-oriented design, analysis, and implementation of a new development version using C++. Using a language with object-oriented features allows for greater modularity of code for both the graphical interface and the data base management system. The use of C++ promotes a software system more resilient to changes in requirements over time and the availability of application objects for other projects in the same domain.

One of the goals facing the VISTA team is to ultimately produce a system that can be run on multiple X workstations. This can be achieved by designing a pure X/Motif system or by porting the GL windows to other platforms by using tools, such as Open GL (released by SGI to other vendors, such as IBM, DEC, Sun, and HP). These options are currently under analysis by the team.

Enhancements to the prototype system will extend its use while the next version is under design and development. Some enhancements have already been planned, such as the incorporation of higher resolution maps for Earth views, additional graphical queries as part of the Build Query section that are more particular to observational events, the addition of mission analysis windows (including detailed views of platforms and their orbits), the addition of special queries and visualization (such as viewing angle from a spacecraft platform), and the addition of graphics and analysis pipelines that handle the visualization of multiple types of data. These and suggestions from prototype reviewers will help the VISTA team drive a more robust data access and handling system.

For more information about the VISTA system and its use, please contact Dr. Edmund Dombrowski (VISTA Coordinator) at 202-404-7831, or send E-mail to dombrowski@bdcv8.nrl.navy.mil. All rights and privileges to the VISTA name and software are copyrighted, 1991, to the Naval Research Laboratory.

Acknowledgments. The authors would like to thank the additional personnel responsible for the coding and development of the VISTA prototype system, namely Sean Tippet (University of Maryland),

Serge Huzarewicz (PRA, Inc.), and Bruce Shetler (PRA, Inc.). We greatly appreciate the editing, reformatting, and repackaging of this paper by BDC's senior technical writer, Ruth Golembiewski. We are grateful for the support and guidance to the VISTA project as provided by Dr. Herb Gursky (NRL Space Science Division Superintendent). Our appreciation also goes out to the Strategic Science Generation Model (SSGM) team at the Naval Research Laboratory for continual support with design and implementation issues.

REFERENCES

- Dombrowski, E.G., R. DeMajistre, J.R. Williams, D.E. Anderson, and W.A. Snyder, The Backgrounds Data Center (BDC) Visual Interface for Space and Terrestrial Analysis (VISTA), paper presented at the Astronomical Data Analysis Software and Systems (ADASS) conference, November 1991, *Astronomical Society of the Pacific Conference Series*, 25, 69-71, 1992a.
- Dombrowski, E.G., R.F. Dorsey, and W.A. Snyder, VISTA: A Visual Interface for Space and Terrestrial Analysis, paper presented at the International Space Year (ISY) / Earth and Space Science Information Systems (ESSIS) Conference, *Proceedings of the ISY/ESSIS*, February 1992b.
- Dombrowski, E.G., B. Dorland, and W.A. Snyder, VISTA software system demonstration at the Astronomical Data Analysis Software and Systems (ADASS) conference III, October 1993.
- Dombrowski, E.G., W.A. Snyder, and H.M. Heckathorn, Metadata Management and the VISTA System, *Proceedings of the Twenty-Seventh Annual Hawaii International Conference on System Sciences (HICSS-27), Information Systems: DSS/Knowledge-Based Systems, III*, 418-427, IEEE Computer Society Press, 1994; presentation at the HICSS-27, January 1994.
- Snyder, W.A., S.L. Berg, R. DeMajistre, E.G. Dombrowski, W.Q. Sadler, and S. Zasadil, The Backgrounds Data Center, *Proceedings of the 4th SDIO Scene Projection Workshop*, December 1991.
- Snyder, W.A., H. Gursky, H.M. Heckathorn, R.L. Lucke, S.L. Berg, E.G. Dombrowski, and R.A. Kessel, The Backgrounds Data Center, paper presented at the Workshop on Advanced Technologies for Planetary Instruments, April 28-30, 1993.

PATHFINDER: Probing ATMOSPHERIC Flows in an INTEGRATED and Distributed ENVIRONMENT

**R.B. WILHELMSON, D.P. WOJTOWICZ, and
C. SHAW**

*National Center for Supercomputing Applications and
Department of Atmospheric Sciences, University of
Illinois, Urbana, IL*

J. HAGEDORN and S. KOCH*

NASA Goddard Space Flight Center, Greenbelt, MD

PATHFINDER is a software effort to create a flexible, modular, collaborative, and distributed environment for studying atmospheric, astrophysical, and other fluid flows in the evolving networked metacomputer environment of the 1990s. It uses existing software, such as HDF, DTM, GEMPAK, AVS, SGI Explorer, and Inventor to provide the researcher with the ability to harness the latest in desktop to teraflop computing. Software modules developed during the project are available in the public domain via anonymous FTP from NCSA. The address is ftp.ncsa.uiuc.edu, and the directory is /SGI/PATHFINDER.

**Currently affiliated with the Department of Marine, Earth, and Atmospheric Sciences, North Carolina State University*

1. INTRODUCTION

During the 1990s, new hardware and software tools are providing researchers with enhanced capabilities for parallel processing, scientific visualization, data management, and distributed applications. Hardware/software advances range from desktop multimedia personal computers to large massively parallel systems. The latter will eventually provide teraflop (10^{12} floating point operations per second) capabilities that will be useful for pursuing answers to "grand challenge" and other important scientific questions. For example, teraflop capability will enable the study of weather changes occurring across the entire United States to be carried out with unprecedented resolution (1 to 2 km horizontal resolution and 30 vertical levels). This is illustrated in Kaufmann and Smarr (1992, p. 18), where each data array or lattice (e.g., temperature, pressure, and wind) in the simulation model has approximately 500 million elements (more than 100 times those in use today). These arrays will be stored in very large memories with maximum configurations well beyond the multiple billion bytes available today. In addition, the handling and storage of data will be expedited by larger disk farms, the use of new storage media (e.g., new DD-2 tapes hold between 25 and 165 gigabytes of data per tape), and high-speed network communications (multiple gigabits per second). Interrogation and visualization of this vast amount of data will be handled by powerful visualization systems that use computer and software resources distributed through the user's working environment.

A term commonly invoked to describe the user's environment consisting of local and remote computers connected by high-speed networks is the "metacomputer." Productive use of this environ-

ment requires software that allows a researcher to distribute and communicate with a collection of processes running on different machines and to store, retrieve, and share information and data. Toward this end, a software effort began several years ago at the National Center for Supercomputing Applications (NCSA) that uses high-performance computing, networking, and data handling capabilities. The aim of the project, known as NCSA PATHFINDER or simply PATHFINDER (Probing ATmospHeric/AsTropHysical Flows in an INtegrated and Distributed EnviRonment), is to create a flexible, modular, collaborative, and distributed environment for studying atmospheric, astrophysical, and other fluid flows [Wilhelmson et al., 1993]. More specifically, this interactive metacomputing environment is being built to:

- Provide standardized access to and centralized control over files, processes, hardware, and other resources distributed across the metacomputer environment
- Script, monitor, or interactively control a sequence of processes that spans multiple machines
- Direct a variety of existing applications to transparently transfer data between each other and remote servers over high-speed networks using standardized data formats
- Manage, access, and relocate large heterogeneous data sets (anything from images to gigabyte data files), along with their descriptive metadata, in a common directory space even though they may physically reside on a number of different platforms
- Browse or search the contents of these data sets and their descriptive information and user-supplied annotations using an interactive multimedia hypertext tool (Mosaic)
- Visualize and compare multiple data streams (such as model and observational) using a diverse collection of visualization software (such as SGI Explorer) and hardware (such as video laser disk, virtual reality, and high definition display)
- Capitalize on existing software capabilities whenever possible

In this paper, some of the components being developed to achieve these objectives are described.

2. GEMVIS

GEMVIS (GEMPAK VISualization) is an outgrowth of the NASA-funded project titled "A Distributed Analysis and Visualization System for Model and Observational Data." This software is a collection of modules that extends some of the capabilities of GEMPAK [desJardins and Petersen, 1989; desJardins et al., 1991] to a three-dimensional graphics environment. GEMPAK (General Meteorological Package) is a stand-alone meteorological software package that has become increasingly popular throughout the atmospheric science community with distributions to at least 55 institutions, including NASA centers, other government research laboratories and organizations, universities and colleges, and industry (available from COSMIC and Unidata). It is used for ingesting and decoding weather data. Extensive diagnostic capabilities are provided for observational and model data. Map transformations are available, along with capabilities for two-dimensional contour plots, streamlines, wind fields, thermodynamic profiles, and isentropic cross sections [desJardins and Petersen, 1989].

To extend these display capabilities to three dimensions and to include animation, portions of GEMPAK that use GEMPAK grid files (analysis and data transformations) have been encapsulated into SGI Explorer Version 2.0 (1991) and AVS Version 4.0 (1992) modules. These modules can be used with other Explorer and AVS modules for three-dimensional viewing and animation. Explorer and AVS are two different implementations of a general visualization environment in which the user "builds" his or her application by connecting modules needed for a specific purpose.

Modules that have been and are being developed will (E refers to SGI Explorer modules and A to AVS modules):

- Extract an arbitrary two-dimensional cross section vertically (this is based on selecting end points of the cross section on a geographical map) (A,E)
- Read GEMPAK grid files and derive all GEMPAK diagnostic functions along with coordinates (A,E)
- Derive GEMPAK grid diagnostics and coordinates based on data sets other than GEMPAK grid files (E)

- Draw a topographic map of a given area and a three-dimensional topographic surface to be registered in any requested GEMPAK map projection (A,E)
- Generate three-dimensional vector displays with controls on arrow style, scale, and coloring (A,E)
- Reformulate the basic internal data format to include metadata (referred to as a grid) permitting the passage of registration, map projection, and annotation information, as well as provide a missing data indicator (A,E)
- Do simple animation of accumulated rendered images (E)
- Display two- and three-dimensional text and annotations (E)
- Provide a number of small utility modules to fill in gaps in the functionality of existing software (E)

Only the first four items above are related specifically to GEMPAK. The others do not require any GEMPAK functions. Furthermore, the display modules do not use any GEMPAK display routines. No provisions have been made to incorporate GEMPAK display routines into the SGI Explorer or AVS framework. An example of GEMVIS output is shown in Figure 1.




Figure 1. A computer screen of a GEMVIS visualization of a potential vorticity surface for a large storm system approaching the east coast of the U.S. The two-dimensional image contains the distribution of potential temperature of the storm in the vertical plane shown. The topography on the surface is also rendered, and State contours are drawn.

3. DISTRIBUTED TOOLS

It is important to have software communications tools for handling communications between computers to use today's distributed computing environment. The Data Transfer Mechanism (DTM) was developed at NCSA for this purpose. It is an application interface that enables inter-process communication by supporting message passing between distributed processes. Several features make it attractive for scientific computing. First, the message protocol is designed to be simple and efficient when transferring scientific data, such as gridded data, raster images, and polygonal data sets. Second, data type conversion is provided for these classes of messages and has been optimized on several supercomputing architectures. These are important considerations when dealing with extremely large data sets.

Modules have been implemented that utilize DTM with SGI Explorer to retrieve data stored on another computer or to pass data from an active simulation on one computer to the Explorer environment. Data are either stored or passed in the Hierarchical Data Format (HDF), which has been adopted as the baseline EOSDIS standard format for science and science-related data. Specifically, it will be used for standard data product generation and for data ingest and distribution. HDF also provides access to NetCDF files. Other formats require translations, many of which can be handled by software developed by other groups.

One example of the use of these modules was demonstrated at the Special Interest Group on GRAPHics (SIGGRAPH) meeting in Chicago during the summer of 1992. Severe thunderstorm phenomena were explored through coupled model initiation, simulation, analysis, and display. The CRAY Y-MP system at NCSA was connected by a T3 (1.54 megabits per second) communications link to an SGI VGX 440 on the conference exhibit floor. The CRAY Y-MP supercomputer handled certain model computations, while the SGI VGX at the exhibit managed the user interface for handling model initiation and parameter changes as well as two- and three-dimensional graphical displays. Using DTM, new model data were analyzed and visualized under user control as they became available from the executing model simulation.

Real-time animations were created by capturing frames on the SGI as they were produced, storing them on laser disk, and then playing them back. An example of a screen display is shown in Figure 2 and another in a recent article by Comerford (1992).

More recently, the CM-2 and the SGI have been coupled. DTM is used to transfer data from memory on the CM-2 during an ongoing simulation to memory on the SGI. For example, PATH-FINDER modules have been used to visualize data from a density current (or thunderstorm outflow) simulation. The display of live model simulations using relatively small simulation grids (for real-time computations) has been carried out successfully during several video teleconferences held between NSF supercomputer centers over dedicated T1 video teleconferencing links.

4. PATHFINDER MODULES

The modules in PATHFINDER, including those referred to above, can be categorized as modules used in data acquisition, data mapping, data analysis, visual idioms, and data presentation. Data acquisition modules can read HDF files or communicate via DTM with executing flow models. HDF files include scientific data, images, and color palettes. The HDF scientific data reader incorporates a user-friendly user interface with support for multiple lattice outputs and vector fields (e.g., flow or vorticity) as well as full n-D support for I/O. The HDF module (ReadDF) also has backward compatibility with older HDF libraries and provides detailed information of file contents (e.g., calculate min, max).

Data mapping modules are available to process the data. The processing includes arithmetic functions and arbitrary vertical slicing in a three-dimensional domain. The modules for arithmetic operations provide examples or building blocks for users to develop modules using their own arithmetic sequences. The vertical slicing module provides a top-view window of the domain. In this separate window, contours from a selected horizontal plane through the data or from a surface map (e.g., map of the U.S. when GEMPAK is used) are used as a reference for selecting two points (by clicking on two different locations) through which a vertical slice is to be displayed. The vertical slice can be an

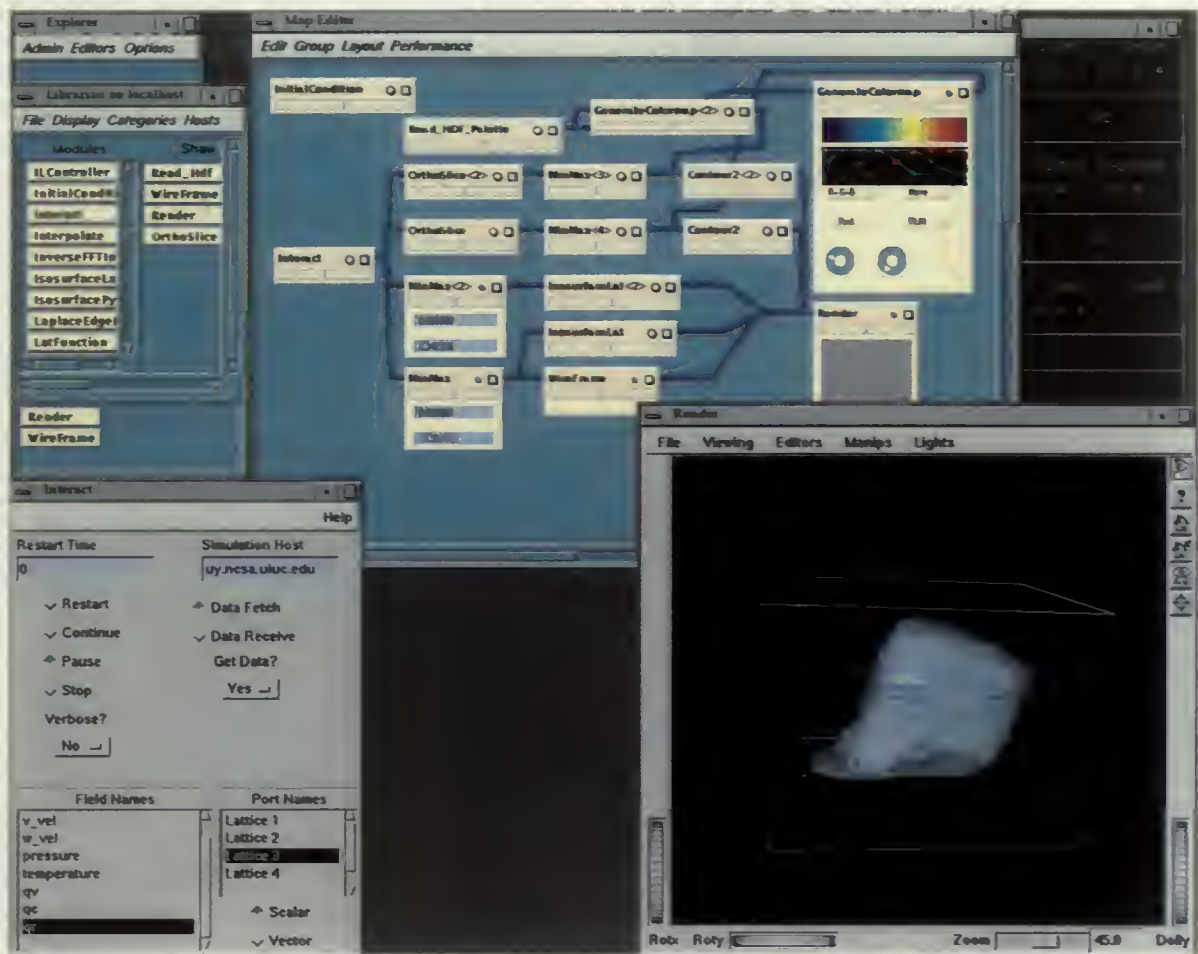


Figure 2. An SGI screen showing a visualization of a growing severe storm along with the associated data flow network, a list of various modules for use in forming a network, an interactive window for starting and stopping the simulation on the Cray Y-MP and matching variables with SGI Explorer lattices, and a window showing min/max values of the data fields as they are computed. The cloud surfaces and horizontal contours are updated automatically as the simulation proceeds.

image, contours, or a set of vectors.

The data analysis modules provide numerical information about the data. They include commonly used functions, such as min/max and trajectory calculations. They also include capabilities for data analysis available in GEMPAK.

Additional and improved visual idioms were developed for PATHFINDER. A new contour module provides the flexibility for selecting contour intervals or the number of contour levels. It has options for showing or hiding zero contour lines and for dimming negative contour lines. The contour data range can be fixed or data dependent. The former is particularly useful for creating contour animations. The new vector module provides options for different vector styles (e.g., line, line with 2/4 cross-line head, line with cross/solid arrow head, and faded line), color mapping based

on other data fields, and different vector scaling. The controls are adjusted according to the range of the input data.

Several annotation modules have also been added to the system. The user can annotate a vertical axis with tick marks and labels, place color maps in the image, and create two-dimensional or three-dimensional titles that incorporate metadata. This was made possible in the Explorer data flow environment by developing a new data type called “grid.” It includes the information in the original lattice data type, as well as descriptive information found, for example, in an HDF file (such as variable name, array dimensions, and time of data).

Particle advection modules have been written to generate points in a user-specified subvolume of the input data domain and to advect all the points

in space for the duration of display interval. Both time-dependent and independent particle trajectories can be computed.

Data presentation modules have been developed for transforming the final image/geometries for use with devices other than the original workstation graphics window. For example, images can be recorded to laser disk and played back on an SGI workstation. In addition, data have been fed in a boom, a virtual reality device that displays stereo images [Sherman, 1992]. The viewpoint changes as the boom is moved by the viewer.

5. ANIMATION TOOL

Animations can be created and played in PATHFINDER, not only using SGI Explorer modules but also using a new tool embedded within PATHFINDER called the Inventor Animator Tool (IATool). Inventor is an object-oriented three-dimensional toolkit for rendering and is used in providing the basic rendering capabilities in SGI Explorer. It utilizes GL commands to produce visualizations from three-dimensional objects, such as isosurfaces, vectors, contours, and motion trajectories. The development of this IATool outside the Explorer environment was motivated by the desire to streamline the creation of animations so that they could be done easily and on a daily basis in a distributed computing environment and based on very large data arrays/sets. These animations involve not only the rendering of objects but also layout, choreography, and annotation. Choreography, including rotation and movement of objects and placement of lighting sources, is important in investigating and communicating important flow features as they evolve over time. Annotation of three-dimensional animations is also important and has lagged behind that available in still images.

The IATool allows a researcher to assemble a detailed and professional-looking animation from three-dimensional objects produced by SGI Explorer or other visualization tools. For example, a researcher might use some Explorer tools (e.g., ReadDF or IsosurfaceLat) to produce isosurfaces of his data. To get the most information from these isosurfaces, they can be animated, be looked at from all sides, and/or integrated with other visual-

ization idioms (e.g., contours or vectors). Using IATool, a sequence in which these tasks are accomplished can be quickly constructed and saved for further inspection or sharing with a colleague.

The IATool allows the specification of both the scene and the animation of the scene in a highly interactive manner from specified geometries. The interface is very similar to that used in the standard Explorer Render module so it should be immediately familiar to Explorer users. Objects are manipulated and positioned with the mouse. They appear exactly as they will in the final product. The position of objects relative to each other or within the frame can be edited along with the object's size and surface properties (e.g., color, shininess, and transparency). Components that effect the whole scene can also be edited, such as lighting and viewpoint (camera position). Furthermore, stereo capabilities are also being implemented for display on a workstation screen or within an immersible virtual reality environment (e.g., the CAVE, a technology environment involving high-resolution rear-projection on multiple walls of a room [Cruz-Neira, 1992]).

In addition to the above features, any of the possible adjustments to the scene can be animated, including the parameters that generate the visualizations. Animation is specified by setting what are called keyframes. Keyframes specify what certain "key" frames of the animation should look like and are most often used in the final production of an animation. The tool uses interpolation methods to generate the frames in between the keyframes. A keyframe is set at the beginning of an animation specifying how everything should appear initially. Another keyframe is then set at some later frame specifying how things should look then. For example, an object might be shown from the front in the first frame and then from behind in a later frame. The viewing position in between keyframes is then automatically interpolated. When all the frames are played, the object would appear to rotate. More complex animations would involve the use of a number of keyframes.

Animations can be replayed directly from within IATool, although the maximum speed is limited by the time it takes to render each frame. For production of a final product, a request would be issued for the tool to render each frame and then

automatically pass the frame to a utility program to generate SGI or MPEG movie files. Other utility programs could send the output to video tape or laser disk (with suitable hardware).

6. THE FUTURE

The evolving electronic and multimedia metacomputing environments of the 1990s will provide researchers with new capabilities that will need to be harnessed through the integration of different computing, communications, data handling, and visualization capabilities. In addition, improved software for working together from remote sites, as well as access to the increasing volume of simulation, observational, and reference data, is needed. High-quality three-dimensional animations will be used routinely and

shared among the research and educational communities. The media for carrying these images will range from high-quality color hard copy for still images and video tapes for animations to multimedia online formats provided through NCSA Mosaic* (see Figure 3) and multimedia electronic mail that combines text, images, movie files (MPEG/Quicktime), and audio. The sharing of information will be available to a large audience as low-cost workstations begin to appear that have

**Mosaic is a cross-platform, hypermedia information retrieval and team collaboration system encompassing existing information servers such as Gopher, WAIS, and WorldWide Web in addition to providing new multimedia services. It is available from NCSA through anonymous FTP. On the Internet use ftp ftp.ncsa.uiuc.edu and logon using anonymous for the name and your local electronic address for the password. Then enter get README.FIRST and follow directions. Enter quit to exit FTP and return to your local host.*

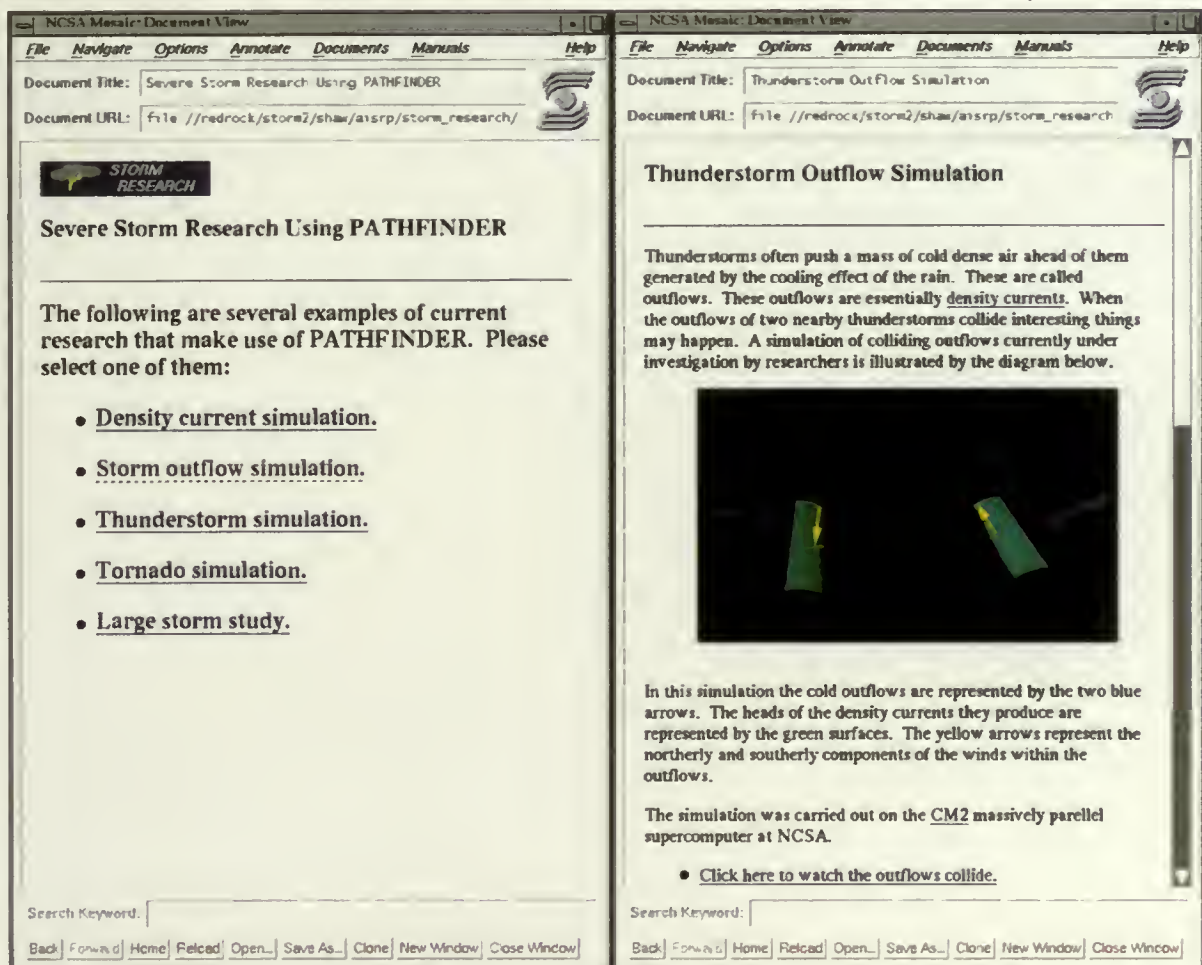


Figure 3. Two pages from a Mosaic document showing examples of severe storm research using PATHFINDER. By clicking on the storm outflow simulation (note dashed underline), the first page of results describing the simulation appears. By clicking at the bottom of the second page, an animation of the outflows as they collide would be shown.

high-quality rendering, stereo, and direct video capabilities. NCSA PATHFINDER and other efforts will strive to meet the evolving needs of the research community.

Acknowledgments. This work has been supported by NASA through the Applied Information Systems Research Program, by NCSA, by NSF through grants ATM 87-00778 and ATM 92-14098, and by Cray Research, Inc. Special thanks to all those who have contributed to the effort, including Matt Arrott, Polly Baker, Brian Jewett, Bruce Lee, Mike McNeill, Guatam Mehrotra, Barbara Mihalas, Bill Sherman, Jeff Terstriep, Jeff Thingvold, and Lou Wicker.

REFERENCES

- Advanced Visual Systems, Inc., *AVS User's Guide*, 1992.
- Comerford, R., Software on the Brink, Special Issue of *IEEE Spectrum on Supercomputers* entitled Teraflops Galore, 34-38, September 1992.
- desJardins, M.L. and R.A. Petersen, *Report on the Final Version of GEMPAK 4*, Preprints, Fifth Conf. on Interactive Information and Processing Systems for Meteorology, Hydrology, and Oceanography, Anaheim, CA, Amer. Meteor. Soc., 205-207, 1989.
- desJardins, M.L., K.F. Brill, and S.S. Schotz, *Use of GEMPAK on UNIX workstations*. Preprints, Seventh International Conf. on Interactive Information and Processing Systems for Meteorology, Hydrology, and Oceanography, New Orleans, LA, Amer. Meteor. Soc., 449-453, 1991.
- Kaufmann III, W.J. and L.L. Smarr, *Supercomputing and the Transformation of Science*, Scientific American Library, New York, 238 pp., 1993.
- SGI, Inc., *IRIS Explorer User's Guide*, 1992.
- Sherman, W.R., Integrating Virtual Environments Into the Dataflow Paradigm, Fourth Eurographics Workshop on ViSC, Workshop Papers, Abingdon, UK, April 1993.
- Wilhelmson, R., S. Koch, M. Arrott, J. Hagedorn, G. Mehrotra, C. Shaw, J. Thingvold, B. Jewett, and L. Wicker, *PATHFINDER-Probing ATmospHeric Flows in an INTERactive and Distributed EnviRonment*, Preprints, Sixth International Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, AMS, 1993.



National Aeronautics and
Space Administration